Planning Guide

hp StorageWorks SAN High Availability

Product Version: FW v06.xx/HAFM SW v08.02.00

Fourth Edition (July 2004)

Part Number: AA–RS2DD–TE/623–000005–001

This guide introduces HP Fibre Channel switching products, storage area networks (SANs), and Fibre Channel technologies. It describes HP StorageWorks directors and edge switches and the *High Availability Fabric Manager (HAFM)* application. It also describes the firmware, backup and restore features, and the graphical user interface delivered with the directors and edge switches and *HAFM* application. Finally, it describes planning for Fibre Channel topologies, physical planning considerations, and configuration planning tasks to ensure taking advantage of director and switch features.

invent

Printed in the U.S.A

SAN High Availability Planning Guide
Fourth Edition (July 2004)
Part Number: AA–RS2DD–TE/623–000005–001

Contents

Figures

Tables

About This Guide

This planning guide provides information to help you plan the acquisition and installation of one or more of the following Hewlett-Packard (HP) products:

- HP StorageWorks Director 2/64
- HP StorageWorks Director 2/140
- HP StorageWorks Edge Switch 2/12
- HP StorageWorks Edge Switch 2/16
- HP StorageWorks Edge Switch 2/24
- HP StorageWorks Edge Switch 2/32
- *High Availability Fabric Manager (HAFM)* application

"About this Guide" topics include:

Overview

This section covers the following topics:

- Intended Audience
- Related Documentation

Intended Audience

This book is intended for use by configuration and installation planners who are experienced with the following:

- System administration.
- Customer engineering.
- Project management.

Related Documentation

For a list of corresponding documentation, see the "Related Documents" section of the Release Notes that came with the product.

For the latest information, documentation, and firmware releases, please visit the HP StorageWorks web site:

http://h18006.www1.hp.com/storage/saninfrastructure.html

For information about Fibre Channel standards, visit the Fibre Channel Industry Association web site, located at http://www.fibrechannel.org.

Conventions

Conventions consist of the following:

- Document Conventions
- Text Symbols
- Equipment Symbols

Document Conventions

This document follows the conventions in Table 1.

Table 1: Document Conventions

Convention	Element
Blue text: Figure 1	Cross-reference links
Bold	Menu items, buttons, and key, tab, and box names
Italics	Text emphasis and document titles in body text
`Monospace font`	User input, commands, code, file and directory names, and system responses (output and messages)
`Monospace, italic font`	Command-line and code variables
Blue underlined sans serif font text (http://www.hp.com)	Web site addresses

Text Symbols

The following symbols may be found in the text of this guide. They have the following meanings:

WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or death.

Caution: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or data.

Tip: Text in a tip provides additional help to readers by providing nonessential or optional techniques, procedures, or shortcuts.

Note: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Equipment Symbols

The following equipment symbols may be found on hardware for which this guide pertains. They have the following meanings:

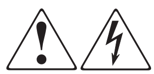

Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts.

WARNING: To reduce the risk of personal injury from electrical shock hazards, do not open this enclosure.

Any RJ-45 receptacle marked with these symbols indicates a network interface connection.

WARNING: To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.

Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

WARNING: To reduce the risk of personal injury from a hot component, allow the surface to cool before touching.

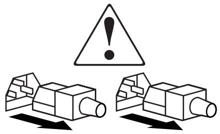

Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

WARNING: To reduce the risk of personal injury from electrical shock, remove all power cords to completely disconnect power from the power supplies and systems.

Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

WARNING: To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

Rack Stability

Rack stability protects personnel and equipment.

WARNING: To reduce the risk of personal injury or damage to the equipment, be sure that:

- The leveling jacks are extended to the floor.
- The full weight of the rack rests on the leveling jacks.
- In single rack installations, the stabilizing feet are attached to the rack.
- In multiple rack installations, the racks are coupled.
- Only one rack component is extended at any time. A rack may become unstable if more than one rack component is extended for any reason.

Getting Help

If you still have a question after reading this guide, contact an HP authorized service provider or access our web site: http://www.hp.com.

HP Technical Support

Telephone numbers for worldwide technical support are listed on the following HP web site: http://www.hp.com/support/. From this web site, select the country of origin.

Note: For continuous quality improvement, calls may be recorded or monitored.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

HP Storage Web Site

The HP web site has the latest information on this product as well as the latest drivers. Access storage at: http://www.hp.com/country/us/eng/prodserv/storage.html. From this web site, select the appropriate product or solution.

HP Authorized Reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518
- In Canada, call 1-800-263-5868
- Elsewhere, see the HP web site for locations and telephone numbers: http://www.hp.com.

Introduction to HP Fibre Channel Products

1

This chapter introduces Hewlett-Packard (HP) Fibre Channel switching products that allow deployment and implementation of a storage area network (SAN) topology in a Fibre Channel Protocol (FCP) or IBM fiber connection (FICON) environment. HP offers several switch alternatives to build a robust and scalable SAN infrastructure that meets the customer's data center requirements. This chapter contains the following:

Product Overview

HP provides three broad classes of Fibre Channel switching products, as follows:

- **Directors** — A director is a high port count, high-bandwidth switch designed with fully redundant, hot-swappable field replaceable units (FRUs) that provide an availability of 99.999% (approximately five minutes of down time per year). HP offers the 64-port StorageWorks Director 2/64 and 140-port StorageWorks Director 2/140.

  The director implements Fibre Channel technology that provides high-performance scalable bandwidth at 2 gigabits per second (Gbps), highly available operation, redundant switched data paths, long transmission distances (up to 50 kilometers at 2 Gbps or 100 kilometers at 1 Gbps), and high device population. Refer to "Directors" on page 22 for detailed information.

- **Edge switches** — An edge switch is a low to medium port count, high-bandwidth switch designed with redundant power supplies and cooling fans that provide an availability of 99.9% (approximately 8.8 hours of down time per year). HP offers the 12-port StorageWorks Edge Switch 2/12, 16-port StorageWorks Edge Switch 2/16, the 24-port StorageWorks Edge Switch 2/24, and the 32-port StorageWorks Edge Switch 2/32 that operate at 2.125 Gbps.

  These switches implement the same high-performance Fibre Channel technology as the director, but with less redundancy, availability, and expense. Refer to "Edge Switches" on page 28 for detailed information.

- **Arbitrated loop switches** — Fibre Channel arbitrated loop (FC-AL) switches are low port count, low-bandwidth products. HP offers related products that act as loop-switching hubs and fabric-attach switches. These switches provide connectivity between attached FC-AL devices, and between FC-AL devices and switched fabric elements. This connectivity allows low-cost or low-bandwidth workgroup (edge) devices to communicate with fabric devices (servers, storage devices, or other peripherals) and ultimately be incorporated into an enterprise SAN environment.

Directors and switches are managed and controlled through a High Availability Fabric Manager (HAFM) appliance, available from HP with the *HAFM*, *Director 2/64 Element Manager*, *Director 2/140 Element Manager*, *Edge Switch 2/16 Element Manager*, *Edge Switch 2/24 Element Manager,* and *Edge Switch 2/32 Element Manager* applications installed. The HAFM appliance is a notebook personal computer (PC) or 1U server that provides a central point of control for up to 48 managed products (directors and switches).

Managed products and the HAFM appliance communicate on a local area network (LAN) through one or more HP-supplied 10/100 Base-T Ethernet hubs. Hubs are daisy-chained as required to provide additional Ethernet connections as more directors or switches are installed on a customer network.

Refer to "Product Management" on page 39 for information about managing products through the HAFM appliance. "Product Management" on page 39 also describes switch management through simple network management protocol (SNMP) workstations, through the Internet using an Embedded Web Server (EWS) interface installed on the product, and through inband (Fibre Channel) application clients.

Directors and switches can be configured to order in an HP-supplied 19-inch equipment rack.

Directors

Directors provide high-performance, dynamic connections between end devices such as servers, mass storage devices, and peripherals in a Fibre Channel switched network. Directors also support mainframe and open systems interconnection (OSI) computing environments and provide data transmission and flow control between device node ports (N_Ports) as dictated by the *Fibre Channel Physical and Signaling Interface* (FC-PH 4.3).

Because of high port count, non-blocking architecture, and FRU redundancy, directors offer high availability and high-performance bandwidth.

Directors should be installed for:

- Backbone implementation for a large-scale enterprise SAN that requires centralized storage management, centralized backup and restore, data protection, and disaster tolerance.
- Mission-critical applications and switched data paths with no downtime tolerance.
- Performance-intense applications that require any-to-any port connectivity at a high bandwidth.

Directors also provide connectivity between servers and devices manufactured by multiple original equipment manufacturers (OEMs). To determine if an OEM product can communicate through connections provided by a director, or if communication restrictions apply, refer to the product publications or contact your HP marketing representative.

Director Performance

Directors provide the following general performance features:

- **High bandwidth** — Each port provides full-duplex serial data transfer at a rate of 2.125 Gbps.
- **High availability** — To ensure an availability of 99.999%, director design provides a redundant configuration of critical components with automatic failure detection and notification.

  Pairs of critical FRUs (logic cards, power supplies, and cooling fans) provide redundancy in case of failure. When an active FRU fails, the backup FRU takes over operation automatically (failover) to maintain director and Fibre Channel link operation. High availability is also provided through concurrent firmware upgrades and spare or unused Fibre Channel ports.

- **Low latency** — The latency is less than 2.5 microseconds between transmission of a frame at a source port to receipt of the frame at the corresponding destination port (with no port contention).
- **Local control** — Actions taking place at a device N_Port seldom affect operation of other ports; therefore, servers need to maintain little or no information about other connected devices in a SAN.
- **Low communication overhead** — Fibre Channel protocol provides efficient use of transmission bandwidth, reduces interlocked handshakes across the communication interface, and efficiently implements low-level error recovery mechanisms. This results in little communication overhead in the protocol and a director bit error rate (BER) less than one bit error per trillion (10^{12}) bits.
- **Multiple topology support** — Directors support both point-to-point and multi-switch fabric topologies and indirectly support arbitrated loop topology.
  - Point-to-point topology provides a single direct connection between two device N_Ports. This topology supports bidirectional transmission between source and destination ports. Through dynamic switching, directors configure different point-to-point transmission paths. In all cases, connected N_Ports use 100% of the available bandwidth.
  - A multi-switch fabric topology provides the ability to connect directors and edge switches through expansion ports (E_Ports) and interswitch links (ISLs) to form a Fibre Channel fabric. Directors receive data from a device, and based on the destination N_Port address, route the data through the fabric (and possibly through multiple switch elements) to the destination device.
  - An arbitrated loop topology connects multiple device node loop ports (NL_Ports) in a loop (or hub) configuration without benefit of a multi-switch fabric. Although directors do not support direct connection of arbitrated loop devices, such devices can communicate with directors through loop switches supplied by HP.
- **Multiple service class support** — The Fibre Channel signaling protocol provides several classes of transmission service that support framing protocol and flow control between ports. Directors support:
  - Class 2 transmission service that provides connectionless multiplexed frame delivery service with acknowledgment. Class 2 service is best suited for mainstream computing applications.
  - Class 3 transmission service that provides connectionless, best-effort multiplexed datagram frame delivery with no acknowledgment. Class 3 service is best suited for mass storage or video applications.

— Class F transmission service that is used by multiple directors to communicate across ISLs to configure, control, and coordinate the behavior of a multi-switch fabric.

Director 2/64

The Director 2/64 is a second-generation, enterprise-class switch that provides switched fabric connectivity for up to 64 Fibre Channel devices. Figure 1 illustrates the front of the director.

Each UPM card provides four 2.125 Gbps Fibre Channel port connections through duplex small form factor pluggable (SFP) fiber-optic transceivers. Shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over single-mode fiber-optic cable. Fiber-optic cables attach to director port transceivers with duplex LC connectors.

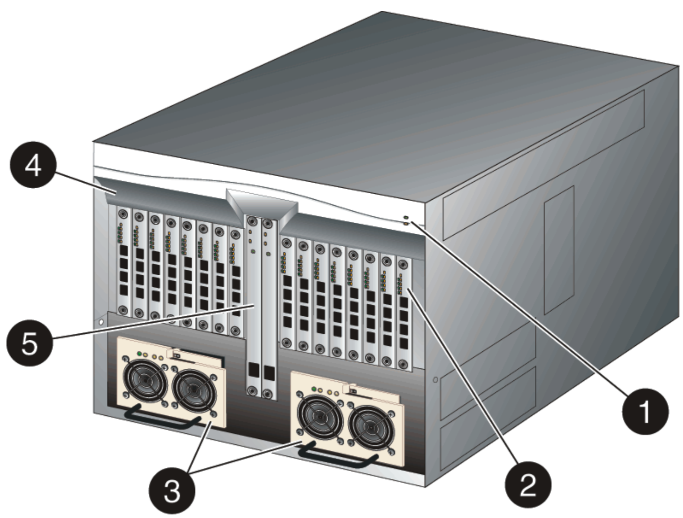

❶ Power and system error light-emitting diodes (LEDs)
❷ Up to 16 universal port module (UPM) cards
❸ Redundant power supplies
❹ Front bezel
❺ Redundant control processor 2 (CTP2) cards

Figure 1: Director 2/64 (front view)

Figure 2 illustrates the rear of the director.

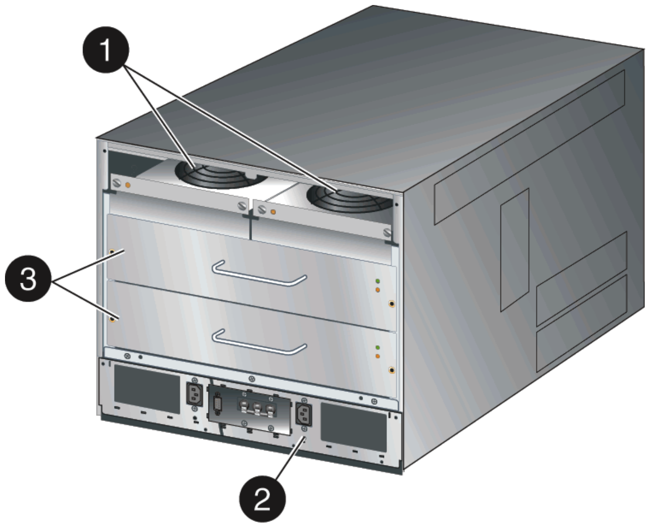

❶ Redundant fan modules
❷ Power module assembly with AC power switch
❸ Redundant serial crossbar (SBAR) assemblies

Figure 2: Director 2/64 (rear view)

The director provides a modular design that enables quick removal and replacement of FRUs. The power module assembly at the rear of the director also provides a 9-pin, D-type subminiature (DSUB) maintenance port for connection to a local terminal or remote terminal. Although the port is typically used by authorized maintenance personnel, operations personnel can use the port to configure director network addresses.

Director 2/140

The Director 2/140 is a third-generation, enterprise-class switch that provides switched fabric connectivity for to up to 140 Fibre Channel devices. Figure 3 illustrates the front of the director.

Each UPM card provides four 2.125 Gbps Fibre Channel port connections through duplex small form factor pluggable (SFP) fiber-optic transceivers. Shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over single-mode fiber-optic cable. Fiber-optic cables attach to director port transceivers with duplex LC connectors.

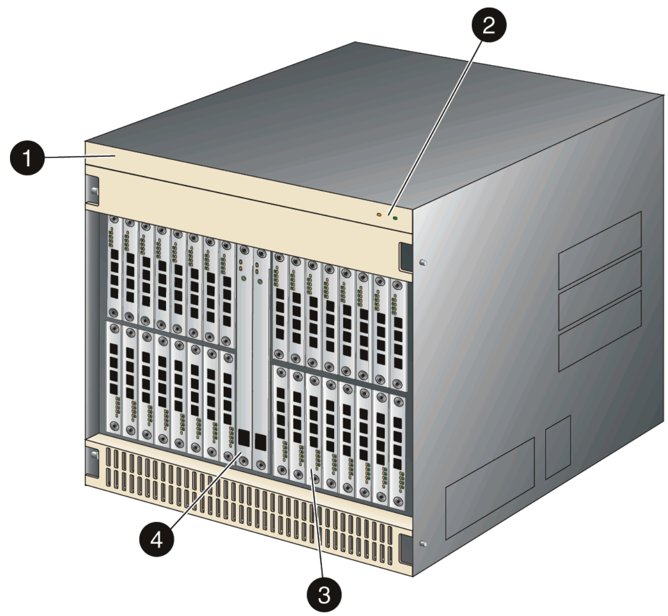

❶	Front bezel	❸	UPM cards (32)
❷	Power and system error LEDs	❹	CTP cards

Figure 3: Director 2/140 (front view)

Figure 4 illustrates the rear of the director.

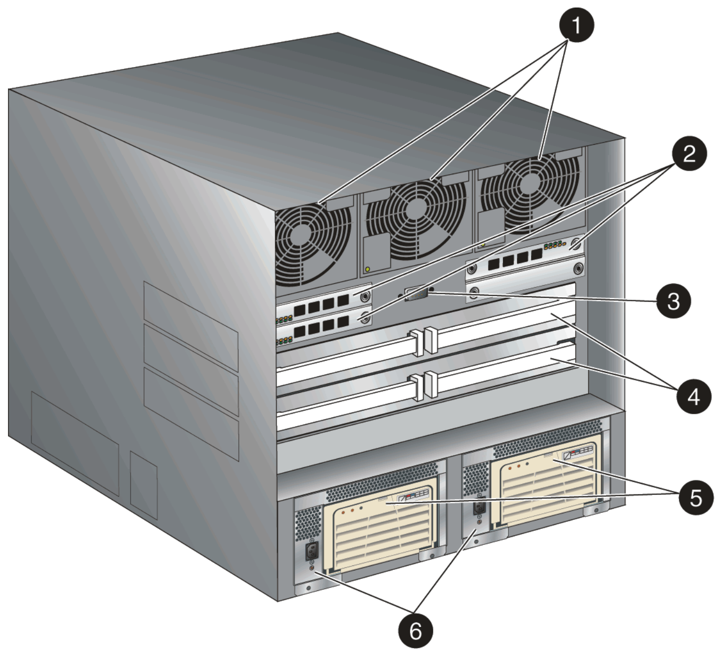

❶ Fan modules
❷ UPM cards (3)
❸ Maintenance port
❹ SBAR assemblies
❺ AC modules
❻ Power supplies

Figure 4: Director 2/140 (rear view)

The director provides a modular design that enables quick removal and replacement of FRUs. The rear of the assembly provides a 9-pin, D-type subminiature (DSUB) maintenance port for connection to a local terminal or remote terminal. Although the port is typically used by authorized maintenance personnel, operations personnel can use the port to configure director network addresses.

Edge Switches

Like directors, edge switches also provide high-performance, dynamic connections between end devices in a Fibre Channel switched network. Edge switches also support mainframe and OSI computing environments.

Through non-blocking architecture and limited FRU redundancy, edge switches also offer high availability and high-performance bandwidth. Although edge switches do not offer the redundancy, availability, or port count of an enterprise-class director, they offer a much lower-cost connectivity option. Edge switches should be installed for:

- Implementation as the principal building block of a small-scale SAN or as a consolidation point for enterprise-class SANs.
- Departmental and workgroup connectivity.
- Applications where distributed storage predominates.

Edge switches also provide connectivity between servers and devices manufactured by multiple OEMs. To determine if an OEM product can communicate through edge switch connections or if communication restrictions apply, refer to the product publications or contact your HP marketing representative.

Edge Switch Performance

Edge switches provide an availability of 99.9% through a redundant configuration of power supplies and cooling fans. When an active FRU (power supply or fan) fails, the backup takes over operation automatically to maintain switch and Fibre Channel link operation. Availability is also provided through concurrent firmware upgrades and spare or unused Fibre Channel ports.

Along with an availability factor of 99.999%, edge switches offer the same general performance features as directors, including high bandwidth, low latency, local control, low communication overhead, multiple topology support, and multiple service class support.

Edge Switch 2/12

The Edge Switch 2/12 provides 2.125 Gbps fabric connectivity for to up to 12 Fibre Channel devices. Figure 5 illustrates the front of the switch.

Shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over single-mode fiber-optic cable. Fiber-optic cables attach to small form factor pluggable transceivers (SFP) with duplex LC connectors. Green and amber status light-emitting diodes (LEDs) are associated with each port.

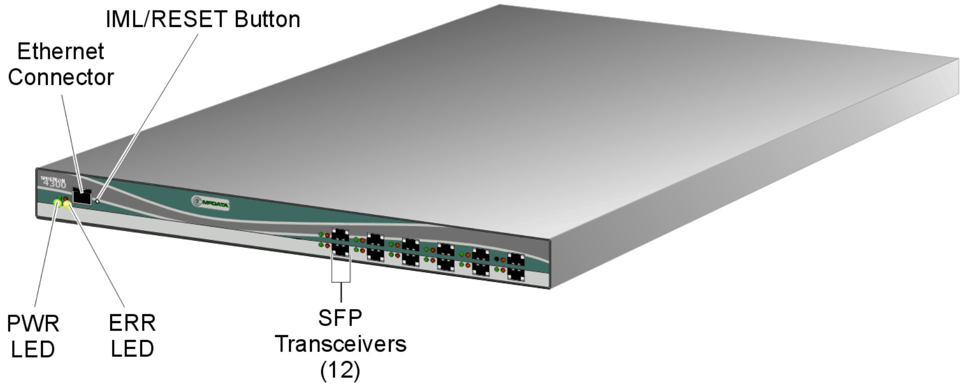

Figure 5: Edge Switch 2/12 (front view)

Figure 6 illustrates the rear of the switch.

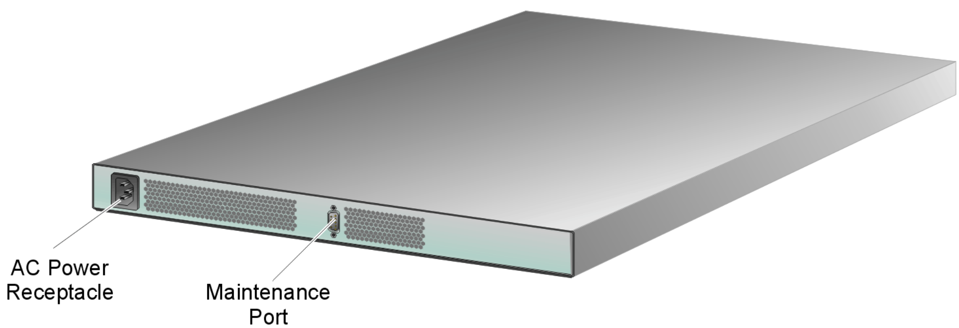

Figure 6: Edge Switch 2/12 (rear view)

Edge Switch 2/16

The Edge Switch 2/16 provides 2.125 Gbps fabric connectivity for to up to 16 Fibre Channel devices. Figure 7 illustrates the front of the switch.

Shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over single-mode fiber-optic cable. Fiber-optic cables attach to switch port transceivers with duplex LC connectors. Green and amber status light-emitting diodes (LEDs) are associated with each port.

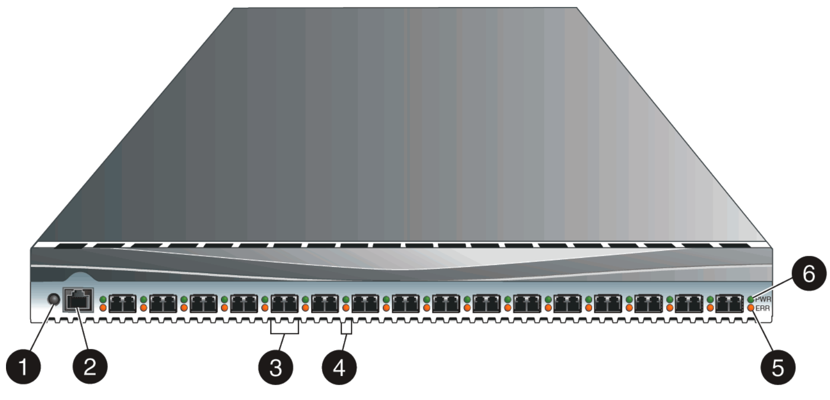

❶	Initial machine load (IML) button	❹	Port LEDs
❷	Ethernet LAN connector	❺	Error LED (amber)
❸	SFP transceivers	❻	Power LED (green)

Figure 7: Edge Switch 2/16 (front view)

Figure 8 illustrates the rear of the switch.

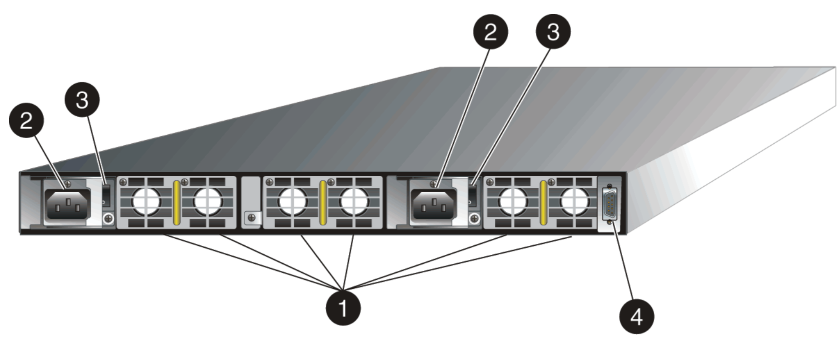

❶	Cooling fans (6)	❸	Power switches on power supplies (2)
❷	Power supplies (2)	❹	RS-232 maintenance port

Figure 8: Edge Switch 2/16 (rear view)

Edge Switch 2/24

The Edge Switch 2/24 provides 2.125 Gbps fabric connectivity for to up to 24 Fibre Channel devices. Figure 9 illustrates the front of the switch.

Shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over single-mode fiber-optic cable. Fiber-optic cables attach to switch port transceivers with duplex LC connectors. Green and amber status light-emitting diodes (LEDs) are associated with each port.

❶	Initial machine load (IML) button	❹	Port LEDs
❷	Ethernet LAN connector	❺	Error LED (amber)
❸	SFP transceivers	❻	Power LED (green)

Figure 9: Edge Switch 2/24 (front view)

Figure 10 illustrates the rear of the switch.

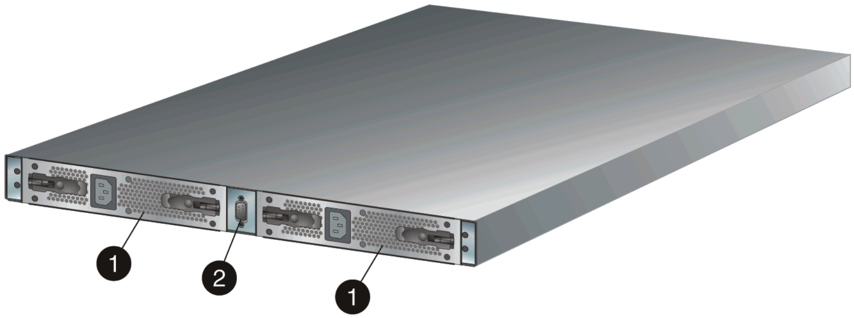

❶ Power supplies with internal cooling fans

❷ Maintenance Port

Figure 10: Edge Switch 2/24 (rear view)

Edge Switch 2/32

The Edge Switch 2/32 provides 2.125 Gbps fabric connectivity for to up to 32 Fibre Channel devices. Figure 11 illustrates the front of the switch.

Shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over single-mode fiber-optic cable. Fiber-optic cables attach to switch port transceivers with duplex LC connectors. Green and amber status LEDs are associated with each port.

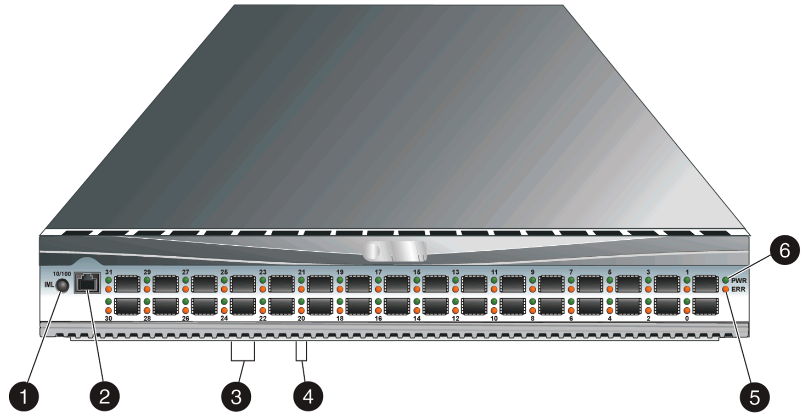

❶ Initial machine load (IML) button
❷ Ethernet LAN connector
❸ SFP transceivers
❹ Port LEDs
❺ Error LED (amber)
❻ Power LED (green)

Figure 11: Edge Switch 2/32 (front view)

Figure 12 illustrates the rear of the switch. The FRUs on the rear panel include two power supplies and four individual cooling fan FRUs.

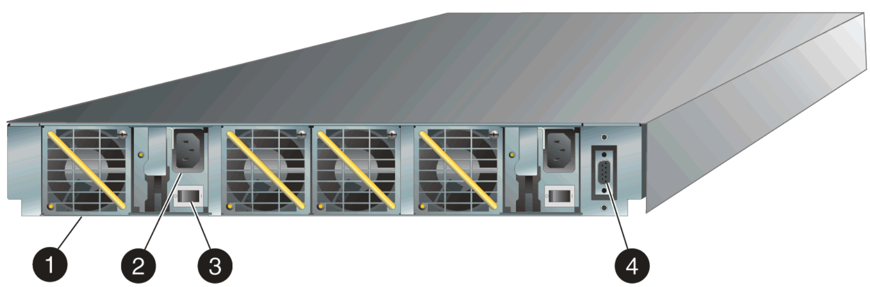

❶ Cooling fans (4)
❷ Power supplies (2)
❸ Power switches on power supplies (2)
❹ RS-232 maintenance port

Figure 12: Edge Switch 2/32 (rear view)

Product Features

In addition to the characteristics and performance features described in this chapter, directors and switches managed by HP also provide a variety of:

- Connectivity Features
- Security Features
- Serviceability Features

Connectivity Features

Directors, switches, and the associated *HAFM* and *Element Manager* applications support the following Fibre Channel connectivity features:

- **Any-to-any connectivity** — Director and switch software configures hardware routing tables for each source port to provide any-to-any port connectivity. Subject to user-defined restrictions such as port blocking and zoning, directors, and switches define the destination port with which a source port is allowed to communicate and provide any-to-any port connectivity. In addition, directors and switches provide connectivity for both FCP and IBM FICON devices.
- **Extended distance support** — Through the use of repeaters, any director or switch port can be configured for extended distance operation. By setting a port's buffer-to-buffer credit (BB_Credit) value to 60, the port can transmit data up to 100 kilometers at 1 Gbps or 50 kilometers at 2 Gbps.
- **Port blocking** — System administrators can block or unblock any director or switch port through the *HAFM* application. Blocking a port prevents an attached device from logging in to the product or communicating with any attached device. A blocked port continuously transmits an offline sequence (OLS).
- **Zoning** — System administrators can partition attached devices into restricted-access zones. Devices in the same zone can recognize and communicate with each other through switched port-to-port connections. Devices in separate zones cannot recognize and communicate with each other.
- **Broadcast and multicast support** — Directors and switches support transmission of a Fibre Channel frame to all attached N_Ports (broadcast) or transmission of a Fibre Channel frame to a user-specified group of attached N_Ports (multi-cast).

- **State change notification** — Directors and switches support a state change notification function that allows attached N_Ports to request notification when other N_Ports change operational state.
- **Port binding** — Directors and switches support a feature that binds an attached Fibre Channel device to a specified port through the device's World Wide Name (WWN).

Security Features

The *HAFM* and *Element Manager* applications offer the following security features:

- **Password protection** — Users must provide a user name and password to log in to the HAFM appliance and access managed directors and switches. Administrators can configure user names and passwords for up to 16 users and can authorize or prohibit specific management permissions for each user.
- **Remote user restrictions** — Remote user access to directors and switches is either disabled or restricted to configured IP addresses.
- **SNMP workstation restrictions** — SNMP workstations can access only management information base (MIB) variables managed by a director or switch SNMP agent. SNMP workstations must belong to SNMP communities configured through the *HAFM* application or EWS interface. If configured, the agent can send authorization failure traps when unauthorized SNMP workstations attempt to access a director or switch.
- **Audit log tracking** — Configuration changes to a director or switch are recorded in an audit log stored on the HAFM appliance, where they are accessible to users for display. Log entries include the date and time of the configuration change, a description of the change, and the source of the change.
- **Port blocking** — System administrators can block or unblock any port to restrict device access to a director or switch.
- **Zoning** — System administrators can create zones that provide director or switch access control to increase network security, differentiate between operating systems, and prevent data loss or corruption. Zoning can be implemented in conjunction with server-level access control and storage device access control.

- **SANtegrity™ Binding** - This feature enhances data security in large and complex SANs that have numerous fabrics and devices provided by multiple OEMs. The feature allows or prohibits director or switch attachment to fabrics (fabric binding) and Fibre Channel device attachment to directors or switches (switch binding).

Serviceability Features

Directors, switches, and the associated *HAFM* and *Element Manager* applications support the following serviceability features:

- LEDs that provide visual indicators of hardware status or malfunctions. LEDs are provided on:
  - Director and switch FRUs.
  - The director front bezel.
  - Switch front panels.
- System alerts, event logs, audit logs, link incident logs, and hardware logs that display director, switch, Ethernet link, and Fibre Channel link status on the HAFM appliance.

  Directors and switches also have threshold alerts and a threshold alert log that notifies users when the transmit (Tx) or receive (Rx) throughput reaches a specified value for configured ports or port types.
- Diagnostic software that performs power-on self tests (POSTs) and port diagnostics (internal loopback and external loopback tests). The software also includes a diagnostic Fibre Channel (FC) wrap test. The FC wrap test applies only when a director or switch is configured to operate in FICON management style.

Note: The Edge Switch 2/12 and Edge Switch 2/24 do not support operation using the FICON management style.

- Automatic notification of significant system events (to support personnel or administrators) through alphanumeric pager, e-mail messages, or the call-home feature.
- An internal modem in the HAFM appliance for HP call-home support.

Note: For directors and switches installed in some legacy environments, call-home notification requires installation of HP Proactive Service software. This service is offered at no additional charge for subsystems covered under an on-site warranty or on-site storage hardware support contract. To register or order Proactive Service software, contact your HP customer service representative.

- An RS-232 maintenance port at the rear of the director or switch (port access is password protected) that enables installation or service personnel to change the product's IP address, subnet mask, and gateway address. The port also allows service personnel to run diagnostics and isolate system problems through a local or remote terminal.
- Redundant FRUs (logic cards, port transceivers, power supplies, and cooling fans) that can be removed or replaced without disrupting director, switch, or Fibre Channel link operation.
- A modular design that enables quick removal and replacement of FRUs without the use of special tools or equipment.
- Concurrent port maintenance. Director UPM cards and switch port transceivers that can be removed, added, or replaced without interrupting other ports or product operation. In addition, fiber-optic cables can be attached to ports without interrupting other ports or product operation.
- Beaconing to assist service personnel in locating a specific port, FRU, director, or switch in a multi-switch environment. When port beaconing is enabled, the amber LED associated with the port flashes. When FRU beaconing is enabled, the amber (service required) LED on the FRU flashes. When unit beaconing is enabled, the system error indicator on the product flashes. Beaconing does not affect port, FRU, director, or switch operation.
- Data collection through the associated *Element Manager* application to help isolate system problems. The data includes a memory dump file and audit, hardware, and engineering logs.
- Status monitoring of redundant FRUs and alternate Fibre Channel data paths to ensure continued director or switch availability in case of failover. The *HAFM* application queries the status of each backup FRU daily. A backup FRU failure is indicated by an illuminated amber LED.
- SNMP management using the Fibre Alliance MIB that runs on the HAFM appliance. Up to 12 authorized management workstations can be configured through the *HAFM* application to receive unsolicited SNMP trap messages. The trap messages indicate operational state changes and failure conditions.

- SNMP management using the Fibre Channel Fabric Element MIB (Version 3.1), Transmission Control Protocol/Internet Protocol (TCP/IP), MIB-II definition (RFC 1213), or a product-specific MIB that runs on each director or switch. Up to six authorized management workstations can be configured through the associated *Element Manager* application to receive unsolicited SNMP trap messages. The trap messages indicate operational state changes and failure conditions.

Product Management

2

This chapter describes management of HP directors and edge switches. The chapter specifically describes:

Product Management Overview

Out-of-band (non-Fibre Channel) management access to HP products is provided through two Ethernet LAN connections to director control processor (CTP) cards or a single connection to a director front panel. The following out-of-band management access methods are provided:

- Management through the *HAFM* application. The *HAFM* application includes the *Director 2/64 Element Manager*, *Edge Switch 2/16 Element Manager*, and *Edge Switch 2/32 Element Manager* applications. This GUI resides on the HAFM appliance and provides a single point of management for all directors and switches. Refer to "Product Software" on page 50 for information about these applications.

  Operators at remote workstations can connect to the HAFM appliance through the local *HAFM* application and associated *Element Manager* applications to manage and monitor directors and switches controlled by the HAFM appliance. A maximum of nine concurrent users (including a local user) can log in to the *HAFM* application. Refer to "Remote User Workstations" on page 141 for information.

Note: Product management through the SAN management or *Element Manager* application is not supported for the Edge Switch 2/12.

- Management using simple network management protocol (SNMP). An SNMP agent is implemented through the *HAFM* application that allows administrators on SNMP management workstations to access product management information using any standard network management tool. Administrators can assign Internet Protocol (IP) addresses and corresponding community names for up to six SNMP workstations functioning as SNMP trap message recipients. Refer to "SNMP Management Workstations" on page 143 for information.
- Management through the Internet using the EWS interface installed on the director or switch. This interface supports configuration, statistics monitoring, and basic operation of the product but does not offer all the capabilities of the corresponding *Element Manager* application. Administrators launch the Web server interface from a remote PC by entering the product's IP address as the Internet uniform resource locator (URL), then entering a user name and password at a login screen. The PC browser then becomes a management console.

- Management through a PC-based Telnet session using the CLI. Any platform that supports Telnet client software can be used.

Figure 13 illustrates an example of out-of-band product management. In the figure, the managed product is a Director 2/64. The customer intranet could be an HP Ethernet hub providing device connectivity.

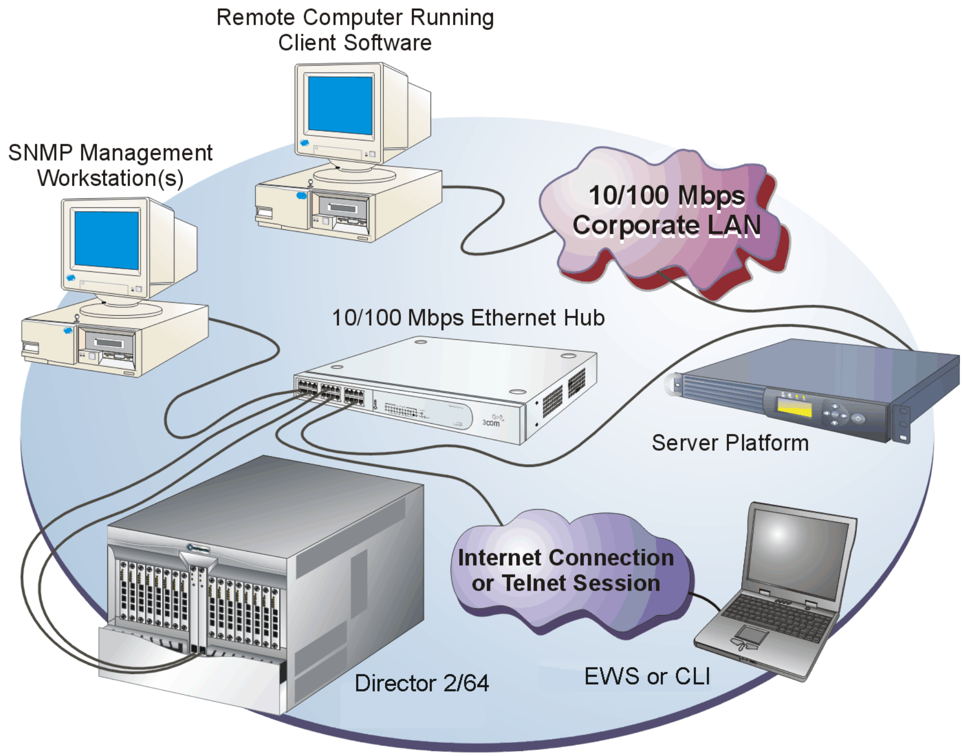

Figure 13: Out-of-band product management

The following inband management access methods are provided as options:

- Management through the product's Open Systems management server (OSMS) that communicates with an application client. The application resides on an open-systems interconnection (OSI) device attached to a director or switch port and communicates using Fibre Channel common transport (FC-CT) protocol. Product operation, port connectivity, zoning, and fabric control are managed through a device-attached console. Refer to "Inband Management Access (Optional)" on page 145 for information.

- Management through the product's Fibre Connection (FICON) management server (FMS) that communicates with one of the following:
  - — IBM® System Automation for OS/390™ (SA OS/390™) operating system resident on a System/390® (S/390) Parallel Enterprise Server™ - Generation 5 or Generation 6.
  - — IBM z/OS® operating system resident on an eServer™ zSeries® 800 (z800), zSeries 900 (z900), or zSeries 990 (z990) processor.

  The server is attached to a director or switch port and communicates through a FICON channel. Control of connectivity and statistical product monitoring are provided through a host-attached console. Refer to "Inband Management Access (Optional)" on page 145 for information.

Figure 14 illustrates inband product management. In the figure, the managed product is a Director 2/64.

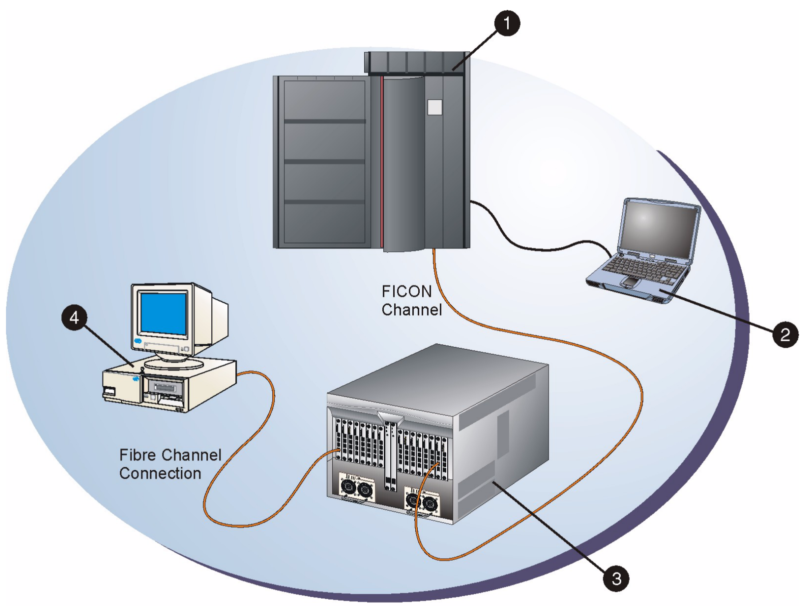

❶ S/390 or zSeries 900 parallel Enterprise server
❷ Host-attached console
❸ Director 2/64
❹ OSI server

Figure 14: Inband product management

HAFM Appliance Description

The HAFM appliance is a one rack unit (1u) high, LAN-accessed, rack-mount unit that provides a central point of control for up to 48 LAN-connected directors or switches. However, note that the maximum number of switches per storage area network (SAN) fabric is different. For the latest supported topology limits, contact your local HP sales representative.

The HAFM appliance is accessed through a LAN-attached PC and standard Web browser. Figure 15 illustrates the HAFM appliance with attached liquid crystal display (LCD) panel.

Figure 15: HAFM appliance

The HAFM appliance is rack-mounted in the HP-supplied FC-512 Fabricenter equipment cabinet. The HAFM appliance or Ethernet access to the *EWS* application is required to install, configure, and manage a director or switch. Although a configured product operates normally without HAFM appliance intervention, an attached HAFM appliance should operate at all times to monitor product operation, log events and configuration changes, and report failures.

The HAFM appliance is dedicated to operation of the *HAFM* and *Element Manager* applications for the following products: Director 2/64, Director 2/140, Edge Switch 2/16, Edge Switch 2/24, and Edge Switch 2/32. These applications provide a GUI and management services and implement Web and other server functions. Refer to "Graphical User Interface" on page 51 for additional information about the applications.

Note: The HAFM appliance and *HAFM* application provide a GUI to monitor and manage multiple HP products and are a dedicated hardware and software solution that should not be used for other tasks. HP tests the *HAFM* application installed on the HAFM appliance but does not compatibility test other third-party software. Modifications to the HAFM appliance hardware or installation of additional software (including patches or service packs) may interfere with normal operation.

United States English is the only language supported by the *HAFM* and *Element Manager* applications.

The HAFM appliance provides two auto-detecting 10/100 Mbps Ethernet LAN connectors (RJ-45 adapters). The first adapter (LAN 1) attaches (optionally) to a public customer intranet to allow access from remote user workstations. The second adapter (LAN 2) attaches to a private LAN segment containing switches or managed HP products.

HAFM Appliance Specifications

The following list summarizes hardware specifications for the HAFM appliance notebook platform. Current platforms may ship with more enhanced hardware, such as a faster processor, additional random-access memory (RAM), or a higher-capacity hard drive or removable disk drive.

- HP OmniBook 6200 PC with color monitor, keyboard, keyboard-mounted trackpad (mouse), and power cord
- 18 gigabyte (GB) or greater internal hard drive
- 160 megabyte (MB) or greater RAM
- Removable DVD/CD-ROM drive
- Removable 100 MB disk drive. This replaces the DVD/CD-ROM drive in the media bay after the initial configuration is complete
- 56K internal modem
- One internal 10/100 Mbps Ethernet adapter with RJ-45 connector (provides public LAN interface to directors and remote clients)

Ethernet Hub

The HAFM appliance and managed directors and switches can be connected through a 10/100 Base-T Ethernet hub. Figure 16 illustrates the 12-port hub. The hub can be ordered from HP and is installed at the top front of the equipment rack.

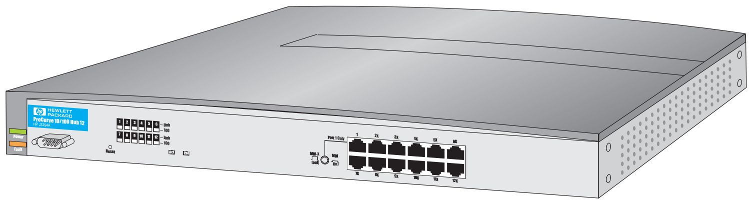

Figure 16: HP Ethernet hub

Remote User Workstations

Operators at remote workstations with the client *HAFM* application installed can connect to the HAFM appliance to manage and monitor all directors and switches controlled by the server. A maximum of 25 concurrent remote users (plus the local HAFM appliance user) can log in to the *HAFM* application. The client application downloads and installs to remote workstations (from the HAFM appliance) using a standard Web browser. The applications operate on platforms that meet the following minimum system requirements:

- Desktop or notebook PC with color monitor, keyboard, and mouse, using an Intel Pentium® III processor with a 1 GHz or greater clock speed and using the Microsoft Windows 2000 Professional (with service pack 3), Windows NT 4.0 (with service pack 6a), or Windows XP (with service pack 1a).
- Unix workstation with color monitor, keyboard, and mouse, using a:
  - — Linux-based system using an Intel Pentium III processor with a 1 GHz or greater clock speed, using the Red Hat® 7.3 or higher operating system.
  - — Hewlett-Packard® PA-RISC® processor with a 400 MHz or greater clock speed, using the HP-UX® 11.0a or higher operating system.
  - — Sun® Microsystems UltraSPARC™ IIi or later processor, using Solaris™ Version 7.0 or higher operating system.
  - — IBM POWER3-II™ microprocessor with a 333 MHz or greater clock speed, using the AIX Version 4.3.3 or higher operating system.
- At least 350 MB available on the internal hard drive.
- 512 MB or greater RAM.

- Video card supporting 256 colors at 800 x 600 pixel resolution.
- Ethernet network adapter.
- Java-enabled Internet browser, such as Microsoft Internet Explorer (Version 4.0 or later) or Netscape Navigator (Version 4.6 or later).

Product Firmware

Director or edge switch firmware provides services that manage and maintain Fibre Channel connections between ports. Although the product hardware transmits Fibre Channel frames between source and destination ports, the firmware maintains routing tables required by the hardware to perform these switching functions. Product firmware also provides functions for system configuration, control, maintenance, and redundancy management, including:

- **System Management Services** — This function configures, controls, and monitors director and switch operation. The subsystem:
  - Centrally manages all configuration and status information.
  - Manages network connections from the HAFM appliance.
  - Implements a simple network management protocol (SNMP) agent to allow access by external SNMP managers using the Fibre Channel Fabric Element management information base (MIB), standard Transmission Control Protocol/Internet Protocol (TCP/IP) MIB-II definition, or product-specific MIB.
- **Fabric Services** — This function supports the fabric controller (login server) and name server. For the director, fabric services also implements a replication manager that synchronizes node port (N_Port) registration databases between redundant control processor (CTP) cards and allows transparent CTP failover.
- **Fibre Port Services** — This function provides a physical driver for hardware components, including:
  - Director 2/64 universal port module (UPM) cards and serial crossbar (SBAR) assemblies.
  - Edge Switch 2/16 and Edge Switch 2/32 fiber-optic ports.
- **Fibre Channel Protocol Services** — This function provides the Fibre Channel transport logic that allows upper layer protocols used by fabric services to communicate with devices attached to fiber-optic ports.
- **Network Services** — This function provides TCP/IP transport layers to access management service subsystems from attached management clients. These clients include the HAFM appliance or an SNMP management station.
- **Application Services** — This function supports all software subsystems for system initialization, logging, tracing, debugging, and communicating with the RS-232 maintenance port.

- **Operating System Services** — This function includes boot and loader software, a command line monitor for engineering fault isolation, a serial maintenance port driver, and other support for the product operating system.
- **Hardware Services (Edge Switch 2/16 and Edge Switch 2/32 Only)** — This function supports the application-specific integrated circuit (ASIC) embedded on the CTP card, provides frame handling for edge switch ports, and provides the application programming interface for light-emitting diodes (LEDs), cooling fans, and power supplies.

Backup and Restore Features

The HAFM appliance provides two backup and restore features. One feature backs up (to the HAFM appliance) or restores the configuration file stored in nonvolatile random-access memory (NVRAM) on a director or switch CTP card. The other feature backs up to a backup drive or restores the entire HAFM data (`HafmData`) directory. The backup and restore features operate as follows:

- **NVRAM configuration** — The NVRAM configuration for any managed director or switch is backed up or restored through the *Element Manager* application. Configuration data (stored in NVRAM on each director or switch) backed up to the HAFM appliance includes:
  - Identification data, such as the director or switch name, description, and location.
  - Port configuration data, such as port names, blocked states, extended distance settings, and link incident (LIN) alerts.
  - Operating parameters, such as buffer-to-buffer credit (BB_credit), error detect timeout value (E_D_TOV), resource allocation timeout value (R_A_TOV), switch priority, and preferred domain ID.
  - Active zoning configuration.
  - SNMP configuration parameters, such as trap recipients, community names, and write authorizations.
- **`HafmData` directory** — Critical information (for all managed products) stored in the `HafmData` directory is backed up or restored using a backup application. The application is configured to automatically back up the contents of the data directory to a removable disk when the HAFM appliance is rebooted or when directory contents change.

  The `HafmData` directory includes:
  - All HAFM configuration data (product definitions, user names, passwords, user rights, nicknames, session options, SNMP trap recipients, e-mail recipients, and Ethernet event notifications).
  - All log files (HAFM logs and individual Element Manager logs).
  - Zoning library (all zone sets and zone definitions).
  - Firmware library.
  - Call-home settings (phone numbers and dialing options).
  - Configuration data for each managed product (stored on the HAFM appliance and in NVRAM on each director or switch).

Product Software

This section describes the *Management Services* and *HAFM* applications. The *HAFM* application includes the *Element Manager* application for each product (Director 2/64, Director 2/140, Edge Switch 2/16, Edge Switch 2/24, and Edge Switch 2/32). The applications provide a GUI and management services for monitoring and controlling directors and switches.

Management Services Application

The *Management Services* application runs on the HAFM appliance and provides management services to the *HAFM* and *Element Manager* applications. It also implements Web and other server functions.

The HAFM appliance is dedicated to the *HAFM* and associated applications and should not be used for other tasks. Loading additional applications or use of the server for other purposes may impact HAFM appliance performance. The *Management Services* application provides the following:

- Session management for one or more HAFM appliance network connections.
- A centralized database repository for configuration files, system logs, firmware upgrades, and other entities.
- Remote support and fault isolation services.
- Establishing and maintaining network connections to managed directors and switches.
- Product configuration management.
- Event and audit logging.
- Alert processing and user notification.
- Initiation of the call-home procedure.

Note: For directors and switches installed in some legacy environments, call-home notification requires installation of HP Proactive Service software. This service is offered at no additional charge for subsystems covered under an on-site warranty or on-site storage hardware support contract. To register or order Proactive Service software, contact your HP customer service representative.

- Network management and file transfer protocol (FTP) processing.

The HAFM appliance also provides hypertext transfer protocol (HTTP) server functionality. Use of this protocol with a standard Web server allows the download of client *HAFM* and *Element Manager* applications from the HAFM appliance to remote workstations. The server is configured to limit the maximum number of concurrent connections to eight.

Graphical User Interface

The HAFM appliance implements the *HAFM* application, along with director and switch-specific *Element Manager* applications, to provide the user interface for operators to control and monitor HP products. These applications can also operate on workstations attached to the customer intranet that function as remote clients.

HAFM Application

The *HAFM* application provides a common Java-based GUI for managed HP products. For more information, see the *HP StorageWorks HA-Fabric Manager User Guide*. The application operates locally on the HAFM appliance or through a network connection from a remote user workstation. The application operates independently from the director or switch managed by the HAFM appliance.

Users can perform the following common product functions:

- Configure new products and their associated network addresses (or product names) to the HAFM appliance for access through the *HAFM* and *Element Manager* applications.
- Display product icons that provide operational status and other information for each managed product.
- Open an instance of an *Element Manager* application to manage and monitor a specific product.
- Display managed fabrics, manage and monitor fabric topologies, manage and monitor zones and zone sets, and show routes (data paths) between end devices attached to a multi-switch fabric.
- Define and configure user names, nicknames, passwords, SNMP agents, and user rights for access to the HAFM appliance, *HAFM* application, and managed products, either locally or from remote user workstations.
- Configure Ethernet events, e-mail notification for system events, and call-home notification for system events.
- Display HAFM audit, event, session, product status, and fabric logs.

HAFM Main Window

The HAFM management application opens automatically when the management server desktop is accessed, and the *HAFM* application main window opens by default.

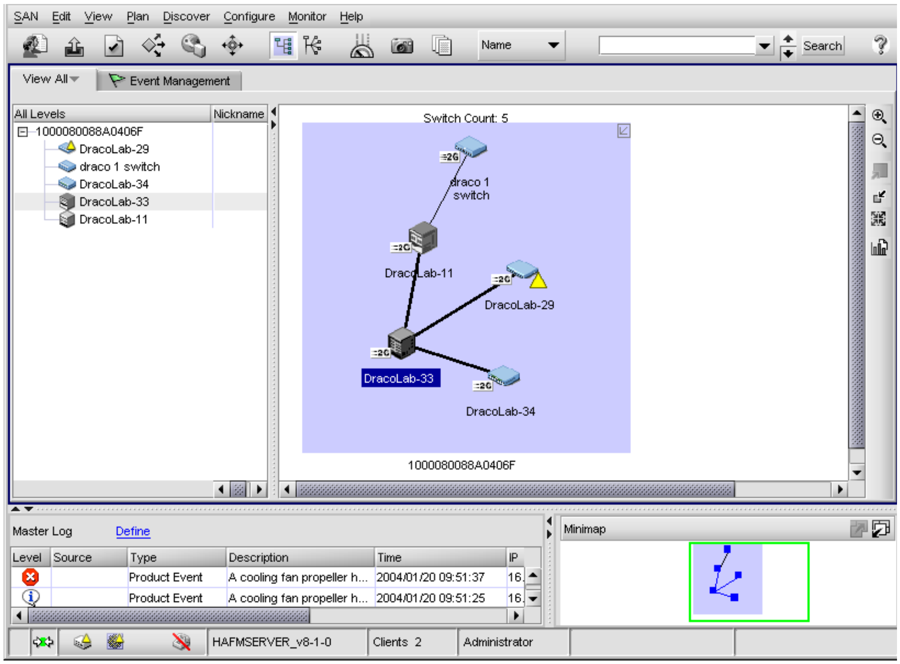

Figure 17: HAFM Main Window

The main window provides the following:

- **Menu bar —** Commands at the top of the window provide drop-down menu selections to perform functions for SAN devices, including editing, viewing, planning, discovery, configuration, and monitoring.
- **Tool bar —** The tool bar (below the menu bar) provides button selections to perform SAN management tasks, including opening a SAN configuration, configuring users, setting up and starting the device discovery process, configuring zoning, displaying a SAN, displaying SAN utilization, and viewing reports.
- **View tab —** Select the **View** tab to display a product list and physical map of the discovered topology.

- **Product list —** When the **View** tab is selected, the product list at the left side of the window displays a list of discovered devices and associated properties.
- **Physical map —** When the **View** tab is selected, the physical map at the right side of the window depicts the SAN topology, discovered devices, and color-coded links.
- **Tool box —** The toolbox at the right side of the window provides button selections to change the discovered topology display, including zoom-in, zoom-out, expand, and collapse functions.
- **Master log —** The master log at the lower left corner of the window displays a list of informational, warning, or fatal events. The log also includes the event source, type, description, time, and IP address of the device generating the event.
- **Utilization legend —** The color-coded utilization legend explains percent utilization for links depicted on the physical map.
- **Minimap —** The minimap at the lower right corner of the window displays the entire SAN topology, and provides an aid to navigate the more detailed physical map.
- **Status bar —** The status bar at the bottom of the window displays connection status, client information, user level, and discovery status.

A label below each icon identifies the managed product. Additional information associated with each icon includes:

- **Data transmission rate —** 2.125 Gbps devices have a 2G label.
- **Attention indicator —** A colored alert symbol adjacent to a product icon indicates the operational status of the product as follows:
  - Absence of an alert symbol indicates the product is fully operational.
  - A yellow triangle indicates a redundant component failure or degraded operational status.
  - A red diamond indicates a critical failure and the product is not operational.
  - A grey square with a yellow exclamation mark indicates the product status is unknown (network connection failure) or the product is offline.

For additional information about the HAFM application, refer to the HA-Fabric Manager User Guide.

Element Manager Application

The *Element Manager* application works in conjunction with the *HAFM* application, and is a Java-based GUI for managing and monitoring multiple directors or switches. The application operates locally on the HAFM appliance, or through a network connection from a remote PC or workstation.

To open an *Element Manager* application, right-click the product icon (Figure 18) at the *HAFM* application's physical map, then select the *Element Manager* option from the pop-up menu. The product icon for an Edge Switch is shown below.

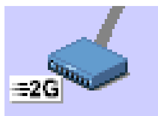

Figure 18: Edge Switch Product Icon

When the *Element Manager* application opens, the last view (tab) accessed by a user opens by default. As an example, the *Hardware View* (Figure 19) for the Edge Switch 2/32 is shown below:

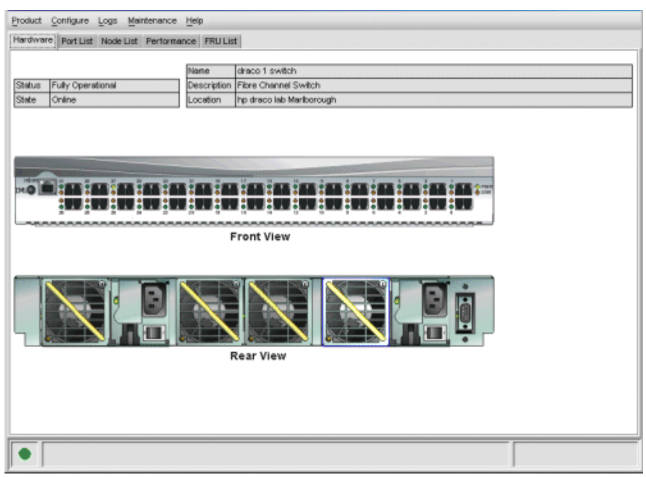

Figure 19: Hardware View

A Director 2/64, Director 2/140, Edge Switch 2/16, Edge Switch 2/24, or Edge Switch 2/32 Status table displays at the top of the window, and a graphical representation of the hardware (front and rear) displays in the center of the window.

The graphical representation of the product emulates the hardware configuration and operational status of the corresponding real product. For example, if a director or switch is fully redundant and fully populated, this configuration is reflected in the **Hardware View**.

Colored symbols display on the graphical FRUs to represent failed or degraded status. The colors and shapes are consistent with status displays on other windows in the *HAFM* and *Element Manager* applications. The light-emitting diodes (LEDs) also highlight to emulate real LED operation.

When the mouse pointer is moved over a FRU in the product graphic, the FRU border highlights in blue and a pop-up identification label displays. Mouse selections (right- or left-click) open dialog boxes or menus that display FRU properties or allow users to perform operations and maintenance tasks.

A menu bar at the top of the **Hardware View** provides **Product**, **Configure**, **Logs**, **Maintenance**, and **Help** options (with associated pop-up menus) that allow users to perform *Element Manager* application tasks.

An Element Manager status bar at the bottom left corner of the **View** window displays colored icons (green circle, yellow triangle, red and yellow diamond, or grey square) that indicate the status of the selected managed product. Messages display, as required, to the right of the colored icons.

Embedded Web Server Interface

With product firmware version 1.2 (or later) installed, administrators and operators with a browser-capable PC and an Internet connection can monitor and manage the director or switch through an EWS interface. The interface provides a GUI similar to the *Element Manager* application and supports product configuration, statistics monitoring, and basic operation.

The EWS interface does not replace nor offer the management capability of the *HAFM* and *Element Manager* applications (for example, the Web server does not support all product maintenance functions). In addition, the EWS interface manages only a single product.

Web server users can perform the following:

- Display the operational status of the director or switch, FRUs, and Fibre Channel ports, and display product operating parameters.
- Configure the product (identification, date and time, operating parameters, and network parameters), ports, SNMP trap message recipients, zones and zone sets, and user rights (administrator and operator).
- Monitor port status, port statistics, and the active zone set, and display the event log and node list.
- Perform product firmware upgrades and port diagnostics, reset ports, enable port beaconing, and set the product online or offline.

The EWS interface can be opened from a standard Web browser running Netscape Navigator 4.6 or higher or Microsoft Internet Explorer 4.0 or higher. At the browser, enter the IP address of the product as the Internet uniform resource locator (URL). When prompted at a login screen, enter a user name and password. The default administrator-level user name is `Administrator`. The default operator-level user name is `Operator`. The default password for both is `password`.

When the interface opens, the default display is the **View** panel (Figure 20). The **View** panel for the Director 2/64 is shown as an example.

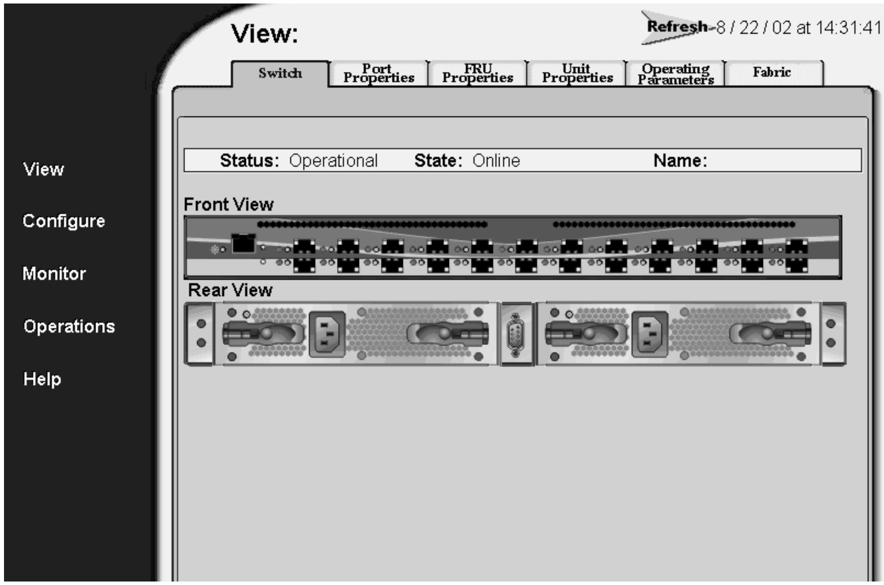

Figure 20: View Panel (Embedded Web Server interface)

Task selection tabs display at the top of the panel, a graphical representation of product hardware (front and rear) displays at the right side of the panel, and menu selections (**View**, **Configure**, **Monitor**, **Operations**, and **Help**) display at the left side of the panel. The task selection tabs allow users to perform director or switch-specific tasks.

The task selection tabs are a function of the menu selected, as follows:

- **View** — On the **View** panel, the **Director** or **Switch** (default), **Port Properties**, **FRU Properties**, **Unit Properties**, **Operating Parameters,** and **Fabric** task selection tabs display.
- **Configure** — On the **Configure** panel, the **Ports** (default), **Switch**, **Management**, **Zoning**, **Security**, and **Performance** task selection tabs display.
- **Monitor** — On the **Monitor** panel, the **Port List** (default), **Port Stats**, **Active Zone Set**, **Log**, and **Node List** task selection tabs display.
- **Operations** — On the **Operations** panel, the **Port Beaconing** (default), **Port Diagnostics**, **Port Reset**, **Online State**, and **Firmware Upgrade** task selection tabs display.
- **Help** — The **Help** selection opens online user documentation that supports the EWS interface.

Command Line Interface

The CLI provides a director and switch management alternative to the *HAFM* application, *Element Manager* application, and EWS user interface. The interface allows users to access application functions by entering commands through a PC-attached telnet session. Any platform that supports telnet client software can be used.

The primary purpose of the CLI is to automate management of several directors or switches using scripts. Although the CLI is designed for use in a host-based scripting environment, basic commands (`config`, `maint`, `perf`, and `show`) can be entered directly at the disk operating system (DOS) window command prompt. The CLI is not an interactive interface; no checking is done for pre-existing conditions, and a user prompt does not display to guide users through tasks.

For additional information, refer to the *HP StorageWorks CLI Reference Guide for Edge Switches and Directors*.

Planning Considerations for Fibre Channel Topologies

3

A storage area network (SAN) is typically defined as a network of shared storage resources that can be allocated throughout a heterogeneous environment. This chapter describes planning considerations for incorporating Hewlett-Packard (HP) switching products into Fibre Channel SAN topologies.

This chapter specifically describes:

Fibre Channel Topologies

The Director 2/64, Director 2/140, Edge Switch 2/12, Edge Switch 2/16, Edge Switch 2/24, and Edge Switch 2/32 support point-to-point and multi-switch fabric topologies. The Edge Switch 2/12 and Edge Switch 2/24 indirectly support arbitrated loop topology. A combination of these topologies (hybrid topology) is also supported.

Related HP switches support switched mode and traditional (shared mode) arbitrated loop topologies and indirectly support a switched fabric topology. These topologies are described as follows:

- **Point-to-point** — This topology provides a single direct connection between two device node ports (N_Ports) and supports bidirectional transmission between the source and destination ports. Through dynamic switching, a director or edge switch configures different point-to-point transmission paths. In all cases, connected N_Ports use 100% of the available bandwidth. For additional information, refer to "Planning for Point-to-Point Connectivity" on page 62.
- **Arbitrated loop** — This topology uses arbitrated loop switches (offered by HP but not described in this publication) to connect multiple device node loop ports (NL_Ports) in an FC-AL or hub configuration without benefit of a multiswitch fabric. The following modes of operation are supported:
  - — Switched mode topology provides a single, logical connection between two device NL_Ports. The switch dynamically configures different logical transmission paths, and in all cases connected NL_Ports have access to 100% of the available bandwidth.
  - — Shared mode arbitrated loop topology connects multiple device NL_Ports in a hub (or star) configuration without benefit of a switched fabric. The switch supports connection of up to 125 arbitrated loop devices and cascaded hubs.

  Although directors and edge switches do not support direct connection of arbitrated loop devices, the loop devices can communicate with fabric elements through an arbitrated loop switch bridge port (B_Port). If peripheral loop devices are expected to communicate with fabric-attached devices, consider installation of a loop switch (with a director or edge switch) to form a fabric-loop hybrid topology. For additional information, refer to "Characteristics of Arbitrated Loop Operation" on page 63, "Planning for Private Arbitrated Loop Connectivity" on page 69, and "Planning for Fabric-Attached Loop Connectivity" on page 75.

- **Multiswitch fabric** — This topology provides the ability to connect directors and edge switches through expansion ports (E_Ports) or interswitch links (ISLs) to form a Fibre Channel fabric. Director or switch elements receive data from a device, and, based on the destination N_Port address, route the data through the fabric (and possibly through multiple switch elements) to the destination device. For additional information, refer to "Planning for Multi-Switch Fabric Support" on page 80, "Planning a Fibre Channel Fabric Topology" on page 99, and "Fabric Topology Design Considerations" on page 109.

Planning for Point-to-Point Connectivity

Point-to-point Fibre Channel topology consists of two device N_Ports communicating by a direct connection through a director or edge switch. The product-operational software provides the ability to configure a dedicated point-to-point connection by binding a director or switch port to a device World Wide Name (WWN).

A dedicated point-to-point connection through a director or switch is simple to implement and should be considered for server-to-storage applications where high performance, high availability, or extended distances are required on a continual basis.

Characteristics of Arbitrated Loop Operation

When implementing Fibre Channel arbitrated loop topology, consideration must be given to switch operating mode, device connectivity, and loop configuration. This section describes the characteristics of arbitrated loop operation, including:

- Shared Mode Versus Switched Mode
- Public Versus Private Devices
- Public Versus Private Loops

Shared Mode Versus Switched Mode

Arbitrated loop switches operate in shared or switched mode as follows:

- **Shared mode** — When set to shared mode, the switch acts as a hub that implements arbitrated loop topology (although the loop has the physical appearance of a star configuration). When a loop circuit is initialized and established, arbitration protocol ensures that only one device attached to a hub port (H_Port) owns the loop at a time.

  The port establishes communication with another device attached to an H_Port (or the B_Port), and half-duplex or full-duplex operation (the default is half duplex) allows the devices to transmit or receive frames at 1.0625 Gbps. During frame transmission between these devices, the full bandwidth of the switch is used and no other H_Ports or devices are available for connection. When frame transmission completes, the loop circuit closes and other devices are able to contend for operation (using standard loop arbitration).

  Shared mode operation is illustrated in Figure 21. Part A shows a server S_2 connected to device D_2 and communicating at 1.0625 Gbps. The B_Port and six remaining H_Ports are inactive. Subsequently, part B shows public device D_1 connected to fabric-attached server S_1, also communicating at 1.0625 Gbps (through the B_Port). The seven remaining H_Ports are inactive.

Note: A Director 2/64 is shown in Figure 21 and other figures as an example. Any HP director or edge switch can be used.

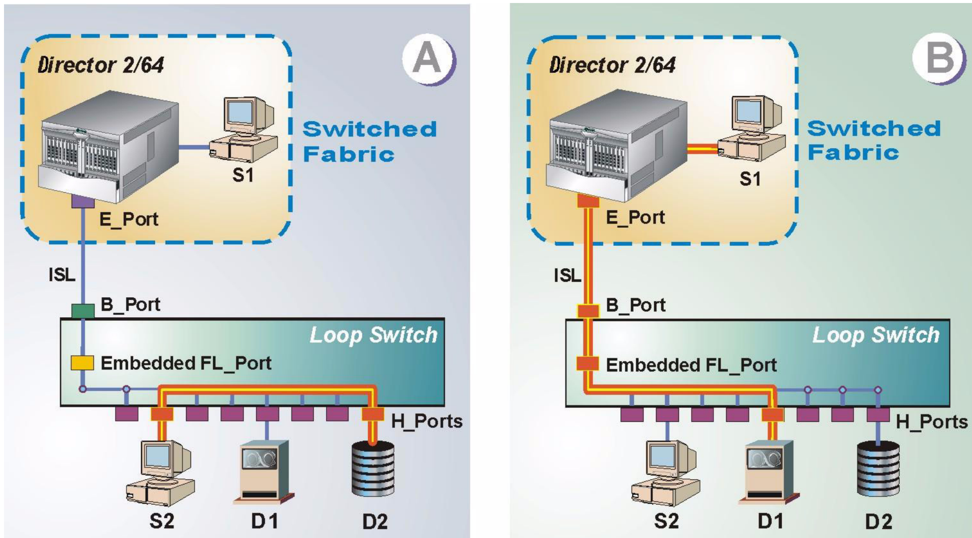

Figure 21: Shared mode operation

- **Switched mode** — When set to switched mode or by default, the switch bypasses full loop arbitration and enables frame transmission through logical connected device pairs. Connections can be established between H_Port pairs, or between an H_Port and fabric loop port (FL_Port). Switched mode also allows independent operation of looplets of devices, each connected through an unmanaged hub, and each attached to a single switch H_Port. Because of opportunistic bandwidth sharing, all looplets or connected device pairs operate half duplex or full duplex at 1.0625 Gbps.

  Switched mode operation is illustrated in Figure 22. Part (A) shows four device transmission pairs using all eight H_Ports (server $\mathbf{S_1}$ to device $\mathbf{D_1}$, server $\mathbf{S_2}$ to device $\mathbf{D_2}$, server $\mathbf{S_3}$ to device $\mathbf{D_3}$, and server $\mathbf{S_4}$ to device $\mathbf{D_4}$). All four transmissions operate independently at 1.0625 Gbps. Subsequently, part (B) shows two device transmission pairs using four H_Ports (server $\mathbf{S_2}$ to device $\mathbf{D_2}$ and server $\mathbf{S_3}$ to device $\mathbf{D_3}$), and public device $\mathbf{D_1}$ connected to fabric-attached server $\mathbf{S_5}$. All four transmissions operate at 1.0625 Gbps.

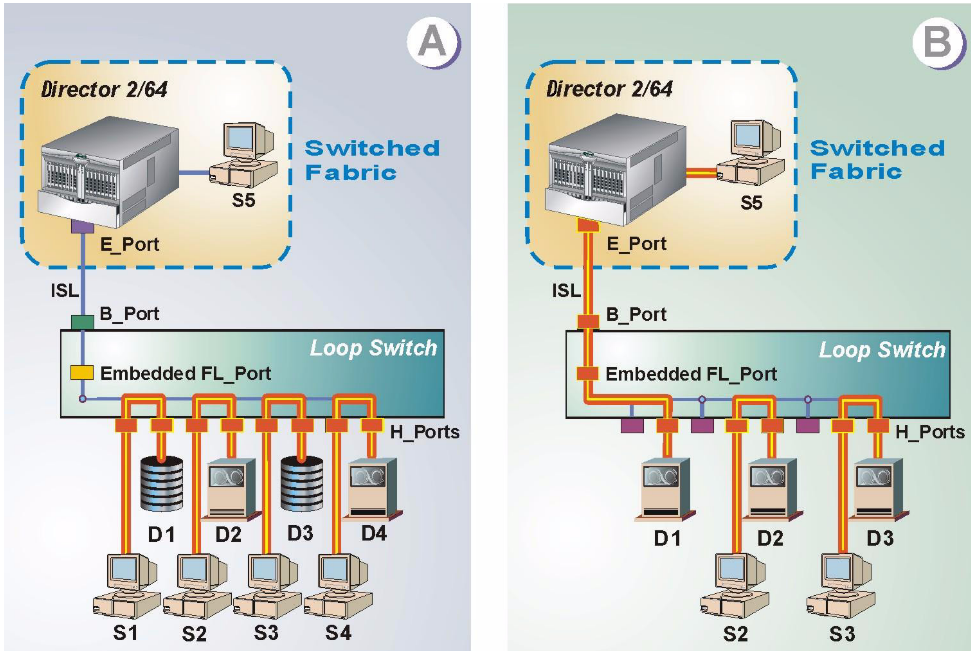

Figure 22: Switched mode operation

Public Versus Private Devices

Arbitrated loop switches support connection of public and private arbitrated loop devices as follows:

- **Public device** — A loop device that can transmit a fabric login (FLOGI) command to the switch, receive acknowledgement from the switch's login server, register with the switch's name server, and communicate with fabric-attached devices is a public device. Public devices communicate with fabric-attached devices through the switch's B_Port connection to a director. As shown in Figure 23, server $\mathbf{S_2}$ is a public loop device that can communicate with fabric-attached device $\mathbf{D_1}$. The switch mode (shared or switched) does not affect device communication.

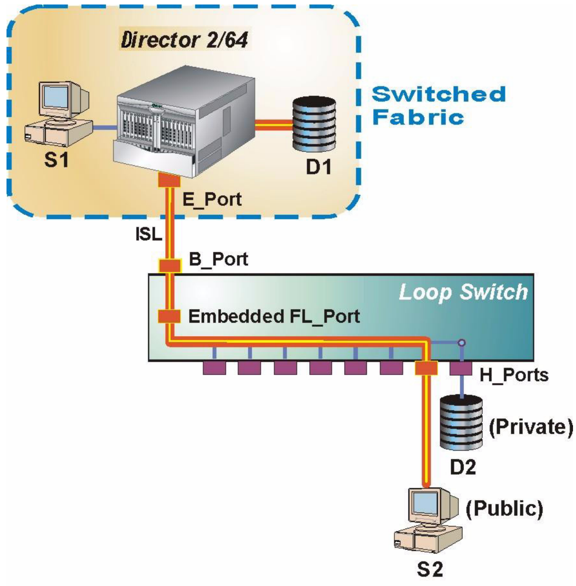

Figure 23: Public device connectivity

Public devices support normal fabric operational requirements, such as fabric busy and reject conditions, frame multiplexing, and frame delivery order.

- **Private device** — A loop device that cannot transmit an FLOGI command to the switch nor communicate with fabric-attached devices is a private device. As shown in Figure 24, device $\mathbf{D_2}$ is a private loop device that cannot communicate with any fabric-attached device. However, device $\mathbf{D_2}$ can communicate with switch-attached server $\mathbf{S_2}$ (using private addressing mode).

  Public and private devices are partitioned into two separate address spaces defined in the Fibre Channel address, and the switch's embedded FL_Port ensures that private address spaces are isolated from a fabric. The switch does not support any other form of Fibre Channel address conversion (spoofing) that would allow private device-to-fabric device communication.

Note: A private device can connect to the switch (loop) while a public device is connected and using the B_Port to communicate with a fabric device.

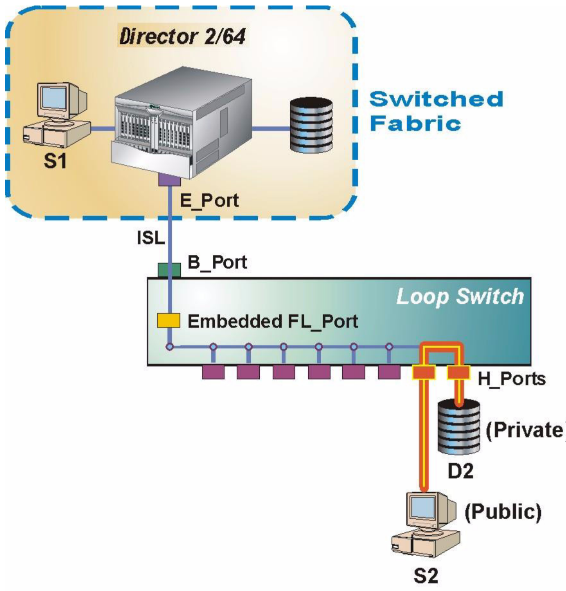

Figure 24: Private device connectivity

Private devices communicate only with other devices on the same arbitrated loop, and interconnected public and private devices can communicate with each other. Such intermixed devices establish operating parameters and loop topology configuration through a port login (PLOGI) command exchange rather than through the switch's name server.

Be aware that public device-to-private device communication may cause problems. For example, it is often critical to separate servers and storage devices with different operating systems, because accidental transfer of information from one to another can delete or corrupt data. Plan to implement security provisions for the switch, such as partitioning attached devices into restricted-access groups (zoning), providing server-level access control (persistent binding), or providing storage-level access control. Refer to "Security Provisions" on page 147 for additional information.

Public Versus Private Loops

Arbitrated loop switches support operation of public and private loops as follows:

- **Public loop** — A public loop is connected to a switched fabric (through the switch B_Port) and the switch has an active embedded FL_Port that is user transparent. All devices attached to the loop can communicate with each

other, and public devices attached to the loop can communicate with fabric-attached devices. FL_Port operation is not affected by the switch operating mode (shared or switched). Public loop connectivity is illustrated in Figure 25.

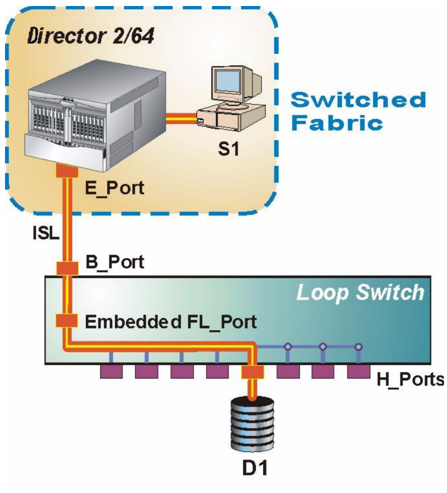

Figure 25: Public loop connectivity

- **Private loop** — A private loop is not connected to a switched fabric, and the switch's embedded E_Port and FL_Port are inactive. All devices attached to the loop can communicate only with each other. Private loop connectivity is illustrated in Figure 26.

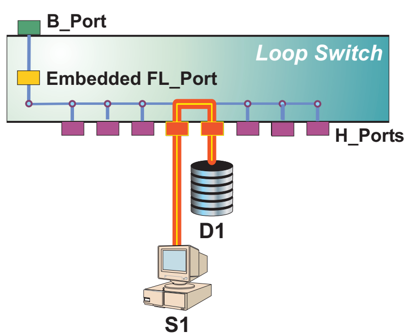

Figure 26: Private loop connectivity

Planning for Private Arbitrated Loop Connectivity

Private arbitrated loop topology supports the clustering of isolated servers and storage subsystems into workgroup or departmental SANs. This topology is well-suited to small-sized and mid-sized configurations where modest connectivity levels and high data transmission speeds are required. The topology also supports low-cost switching and connectivity in environments where the per-port cost of a director is prohibitive.

Private arbitrated loop topology:

- Supports the connection of up to 125 node (device) ports.
- Reduces connection costs by distributing the routing function through each loop port (loop functionality is a small addition to normal Fibre Channel port functionality).
- Provides a fully blocking architecture that allows a single connection between any pair of loop ports at any time. Connections between a third loop port and busy ports are blocked until communication between the first connection pair ends.

Shared Mode Operation

When set to shared mode, a loop switch implements standard Fibre Channel arbitrated loop topology and distributes the frame routing function through each loop port. Shared mode operation and its simplified logical equivalent are illustrated in Figure 27.

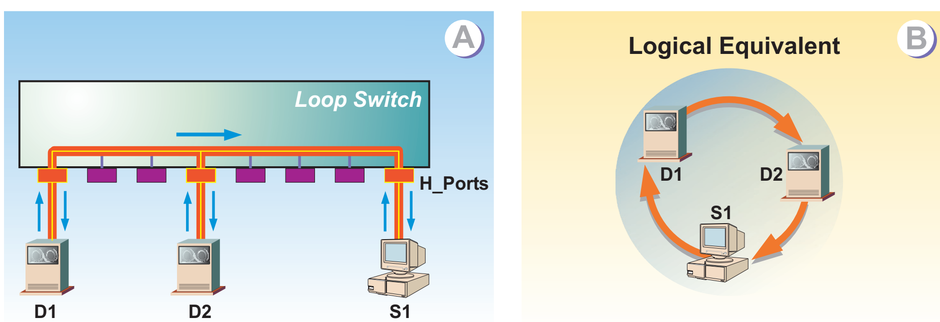

Figure 27: Shared Mode operation and logical equivalent

Part A of Figure 27 shows device D_1 connected to server S_1 through a pair of H_Ports and communicating at 1.0625 Gbps. Although the remaining switch H_Ports (six ports) and device D_2 are unavailable for connection, frame traffic

between device D_1 and server S_1 travels through a loop that consists of all eight H_Ports, device D_1, device D_2, and server S_1. Each H_Port not participating in the communication pair and the NL_Port on device D_2 provide a repeater function that allows frames to pass around the loop at the full switch bandwidth.

Part B of Figure 27 shows the logical equivalent of this arbitrated loop. When frame transmission between device D_1 and server S_1 completes, the loop circuit closes and other ports attached to initiating devices arbitrate for loop access. When operating in shared mode, the switch is a serially reusable resource that provides service access to all ports on the loop. Access is granted by successful arbitration. When arbitration is won by a device, the loop is busy and other arbitrating devices must wait for loop access.

Device attachment and loop construction are not limited to the eight switch H_Ports. Through the use of cascaded unmanaged hubs, the Fibre Channel architectural limit of 125 FC-AL devices can attach to the switch. For example, Figure 28 shows a private loop composed of a loop switch, 20 FC-AL devices, and two unmanaged hubs.

Hubs are cascaded through H_Port-to-H_Port connections (one port per switch or hub). Server S_1 communicates with device D_1 through a loop that includes H_Ports on all three hubs and NL_Ports on the remaining 18 devices.

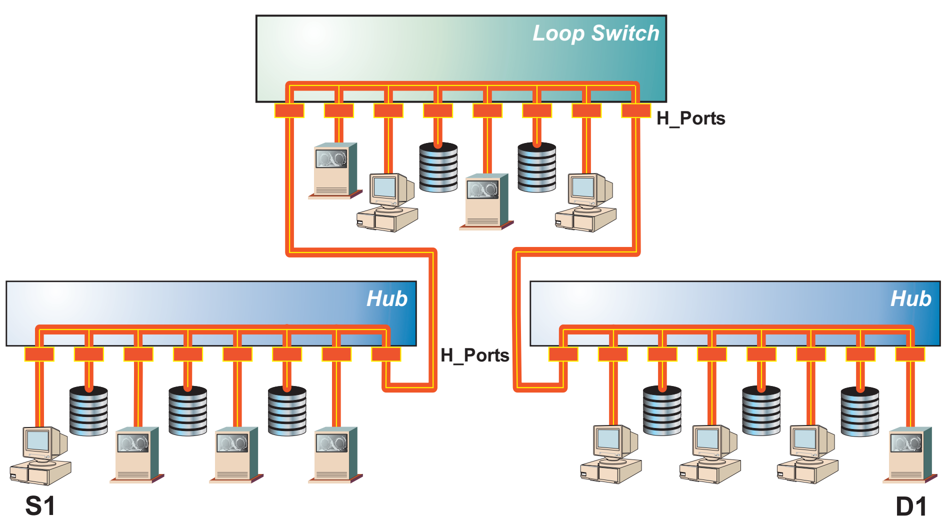

Figure 28: 20-Device private arbitrated loop

Although connection of additional devices to a loop does not impact switch bandwidth (1.0625 Gbps), it does adversely impact overall loop performance because part of the bandwidth is dedicated to overhead instead of information transmission.

Loop performance is a complex function of several factors, including the:

- **Loop round-trip time** — The time required for a frame to travel completely around a loop is a function of the propagation delay associated with each H_Port and NL_Port and the time required to travel through the fiber-optic or copper transmission medium. The addition of ports (through cascaded hubs), devices, and cabling increases the round-trip time.
- **Number of loop tenancies** — Each cycle of device arbitration, loop opening, frame transmission, frame reception, and loop closing is called a loop tenancy. A Fibre Channel operation, such as a small computer system interface (SCSI) write command, may require several tenancies to complete. Because significant overhead is associated with establishing and ending each loop tenancy, an increase in tenancies decreases loop performance. To decrease the number of loop tenancies, plan to limit the number of arbitration-initiating devices installed on the loop.
- **Service rate** — The loop service rate is the number of H_Port service requests the arbitrated loop can process in a time period and is defined as *one* divided by the *average loop tenancy duration*. Long-duration loop tenancies decrease the loop service rate because select devices monopolize the loop. High-bandwidth storage devices that can rapidly process input/output (I/O) requests typically cause long-duration loop tenancies. Plan to limit the number of such devices installed on the loop.
- **Loop utilization** — Loop utilization is the term that describes how busy an arbitrated loop is and is defined as the *request rate* divided by the *service rate*. The request rate is the rate at which devices arbitrate for access to the loop and is a function of applications using the loop, not the loop itself. As the request rate increases due to additional devices being added to the loop, the probability of contention for loop access and arbitration wait time increases. In fact, as loop utilization increases, arbitration wait time increases nonlinearly.

Shared mode operation does not fully use the switch's capabilities and should be used only when connecting legacy FC-AL devices that do not support switched mode operation.

Although the architectural limit of a Fibre Channel arbitrated loop is 125 devices, 32 or fewer devices should be attached to the switch to avoid adversely impacting loop performance. In particular, avoid attaching an excess number of servers or high-bandwidth storage devices.

Switched Mode Operation

When set to switched mode (default setting), a loop switch enables frame transmission through multiple point-to-point connected pairs. Switched mode operation and its simplified logical equivalent are illustrated in Figure 29.

Part A of Figure 29 shows server S_1 connected to device D_1 through a switched pair of H_Ports, communicating at 1.0625 Gbps. Server S_2 is connected to device D_2 through a second switched pair of H_Ports, also communicating at 1.0625 Gbps. Because of opportunistic bandwidth sharing, the two switched pairs effectively increase the switch bandwidth to 2.125 Gbps. The remaining switch H_Ports (four ports) are available for switched connection to each other but cannot communicate with servers S_1 and S_2 or devices D_1 and D_2. Part (B) of Figure 29 shows the logical equivalent of this arbitrated loop.

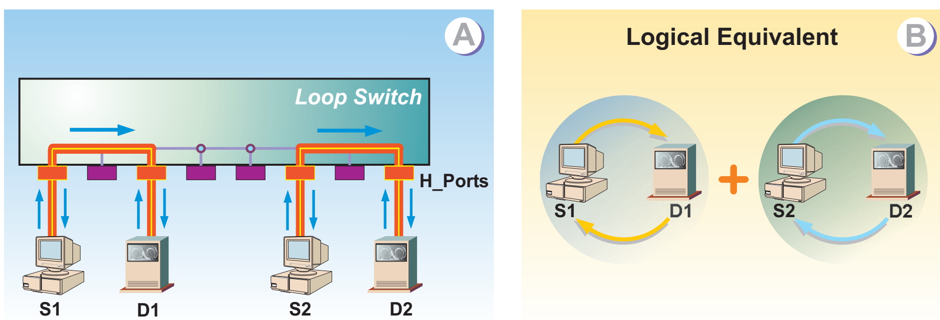

Figure 29: Switched mode operation and logical equivalent

Switched mode also allows independent operation of looplets of devices, each connected through an unmanaged hub and each attached to a single switch H_Port. Figure 30 shows 8 hubs, each connected to a switch H_Port and each connected to a pair of devices (16 devices total). Each device pair forms a looplet that communicates through a hub and connecting H_Port, and because of opportunistic bandwidth sharing the looplets effectively increase the switch bandwidth to 8.5 Gbps.

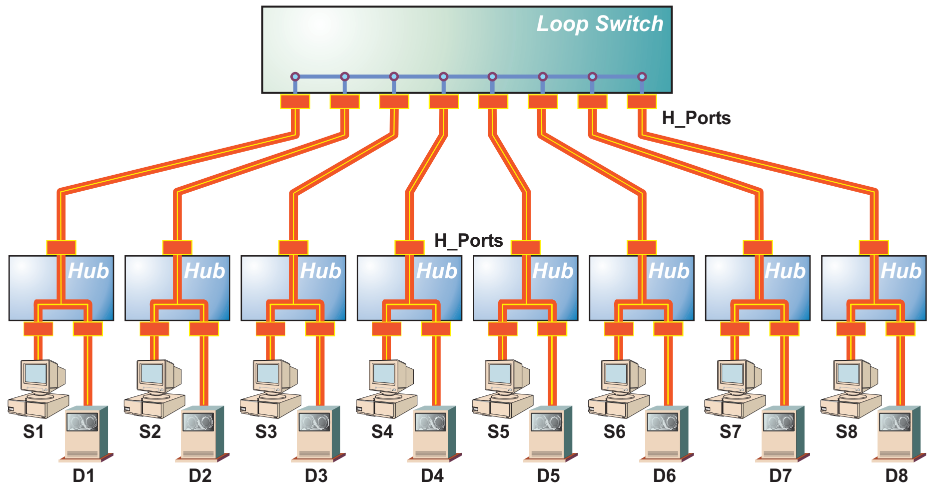

Figure 30: Switched mode operation with eight independent looplets

When communication within two or more looplets ceases, a device attached to one looplet can be switched to communicate with a device attached to another looplet.

Downstream devices in a looplet are attached to an unmanaged hub; therefore, each looplet is operating in normal FC-AL loop mode. However, each looplet attaches to a switched H_Port and ideally should support only connections and operations for up to 32 FC-AL devices. Therefore, an arbitrated loop can be constructed that supports the architectural 125-device limit.

Switched mode operation provides the ability to design a complex and high-performance SAN for the department or workgroup.

Consider the following when planning such a SAN:

- Connect loop switch H_Ports to multiple unmanaged hubs to provide additional FC-AL device connectivity in the form of looplets. Cascade the unmanaged hubs if more than eight hubs are necessary for the configuration.
- Attach devices that frequently communicate with each other to the same looplet to take advantage of opportunistic bandwidth sharing (communication predominately stays within the loop). Switched connections through the loop switch allow connectivity as necessary to devices attached to other looplets.
- Each looplet acts as a normal FC-AL loop. Spread multiple servers and high bandwidth storage devices across several looplets to avoid performance problems associated with a single looplet.

- Consider the data traffic capacity of the department or workgroup (normal and peak load) as part of the switch planning and installation process.

  Such capacity planning:

  — Ensures loop traffic is distributed and balanced across servers and storage devices.

  — Identifies traffic bottlenecks and provides for alternate connectivity solutions if required.

  — Assists in calculating scalability to satisfy nondisruptive growth requirements or eventual connection to a Fibre Channel switched fabric.

  Capacity planning is a dynamic activity that must be performed when new devices, applications, or users are added to the department or workgroup loop configuration.

Planning for Fabric-Attached Loop Connectivity

Public arbitrated loop topology supports the connection of workgroup or departmental FC-AL devices to a switched fabric through a loop switch B_Port. This topology is well suited for:

- Providing connectivity between a workgroup or departmental SAN and a switched fabric, thus implementing connectivity of FC-AL devices to fabric devices at the core of the enterprise.
- Consolidating low-cost Windows NT or UNIX server connections and providing access to fabric-attached storage devices.
- Consolidating FC-AL tape device connections and providing access to fabric-attached servers.

Connecting a SAN to a Switched Fabric

Arbitrated loop switches provide a B_Port that dynamically connects FC-AL devices to directors or edge switches participating in a Fibre Channel fabric. This function allows multiple low-cost or low-bandwidth departmental or workgroup devices to communicate with fabric-attached devices through a high-bandwidth link and provides connectivity as required to an enterprise SAN environment. This approach provides:

- Cost-effective FC-AL device connectivity to a switched fabric. The B_Port provides fabric connectivity without incurring true switched fabric costs. However, the switch does not provide the same simultaneous connection and bandwidth capabilities provided by a Fibre Channel director or switch.
- Improved access and sharing of data and computing resources throughout an organization by connecting isolated departmental or workgroup devices to the core data center. Fabric-to-loop connectivity ensures that edge servers have access to enterprise storage and edge peripherals have access to enterprise computing resources.
- Improved resource manageability. Distributed resources are consolidated and managed through Fibre Channel connectivity instead of physical relocation. One High Availability Fabric Manager (HAFM) appliance manages the operation and connectivity of multiple directors, edge switches, fabric-attached devices, arbitrated loop switches, and FC-AL devices.
- Improved security of business applications and data. Directors, edge switches, and loop switches allow fabric-attached and FC-AL devices to be partitioned into restricted-access zones to limit unauthorized access. Refer to “Zoning” on page 154 for more information.

The switch B_Port provides a single 1.0625 Gbps ISL to an E_Port on a director or edge switch. Direct ISL connectivity between loop switches (with or without a redundant B_Port connection to a director or edge switch) is generally not supported. However, a director or edge switch does support the connection of multiple, independent switches. Figure 31 shows a configuration of two loop switches attached to a director.

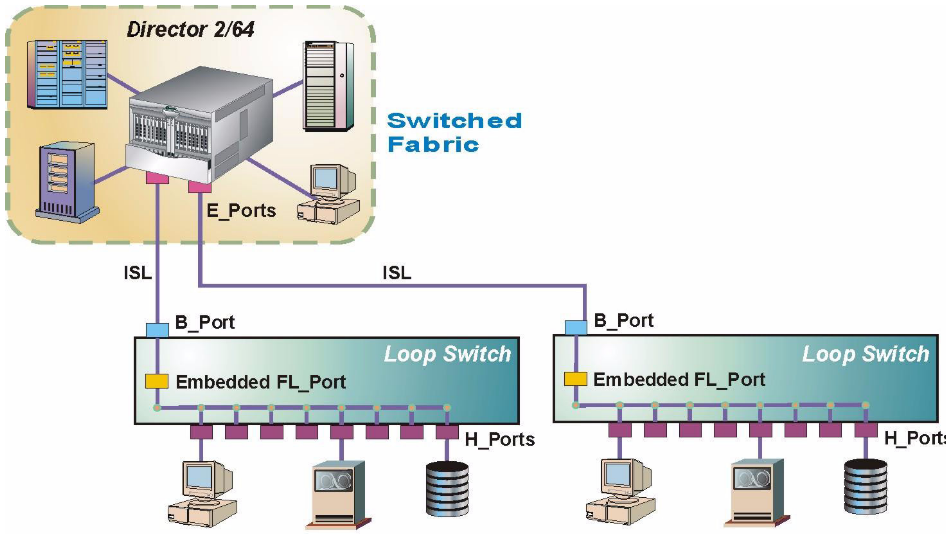

Figure 31: Arbitrated loop to switched fabric connectivity

Consider the following when planning arbitrated loop-to-switched fabric connectivity and incorporating FC-AL devices into the enterprise SAN environment:

- B_Port traffic is routed through a user-transparent FL_Port that is embedded on the switch’s control processor (CTP) card. Switch mode (shared or switched) has no impact on B_Port operation. However, because all switch-attached FC-AL devices must arbitrate for access to the embedded FL_Port, loop performance issues (loop round-trip time, number of loop tenancies, service rate, and loop utilization) must be evaluated.
- Although the B_Port connection (ISL) between the director and switch is a 1.0625 Gbps serial connection, burst transmissions from multiple FC-AL devices are multiplexed and buffered (the link BB_Credit value is eight) and may coexist in the link. Therefore, the sum of the bandwidths of all devices contending for B_Port access should not exceed 1.0625 Gbps. Exceeding the total bandwidth may result in degraded performance, as shown in Figure 32.

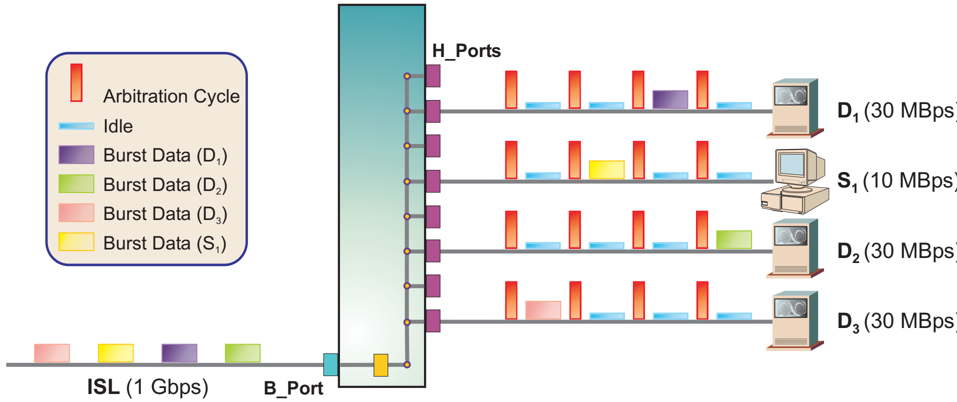

Figure 32: ISL bandwidth limitation

Three 30-megabyte per second (MBps) tape drives and one 10-MBps server are attached to a switch. Each device uses only a portion of the bandwidth of its respective H_Port connection. Each H_Port connection illustrates four arbitration cycles (each cycle won by a single device), one cycle of burst data transmission, and three idle cycles. The B_Port connection to a director (ISL) transmits multiplexed burst data from all four devices, for which the summed bandwidths equal the 1.0625 Gbps capacity of the link. Therefore, connection of additional devices to the switch adversely impacts B_Port performance.

Server Consolidation

Providing fabric connectivity for multiple low-bandwidth servers (Windows NT or UNIX-based) by attaching them individually to an expensive Fibre Channel director is not a cost-effective solution. A practical solution is to consolidate the servers on an inexpensive loop switch, and then connect the switch to a single director or edge switch E_Port.

Figure 33 illustrates the consolidation of ten servers (using two unmanaged hubs) through one B_Port connection to a director. Each server has a 10-MBps bandwidth, therefore, the sum of the bandwidths of all consolidated servers equals the B_Port bandwidth of 1.0625 Gbps. Connecting another server to the switch would exceed the B_Port capability and adversely impact director-to-switch link performance. Other devices (such as tape drives) should not be connected to a switch used for server consolidation.

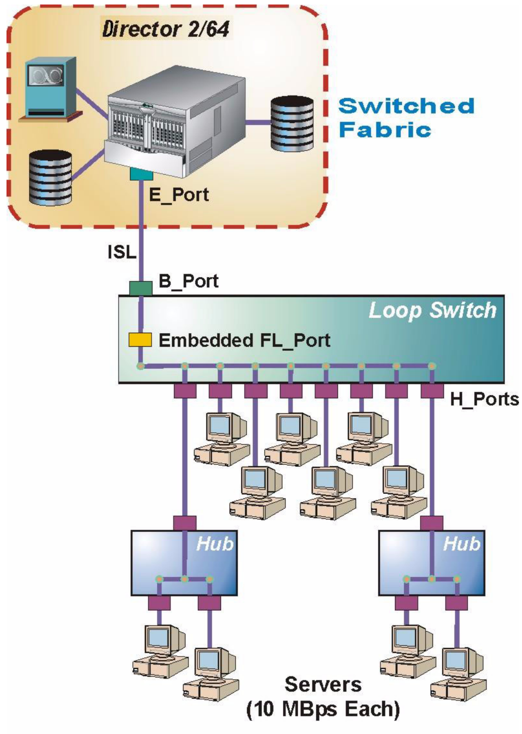

Figure 33: Server consolidation

Tape Device Consolidation

Providing fabric connectivity for multiple FC-AL tape drives by attaching them individually to a Fibre Channel director is likewise not a cost-effective solution. A practical solution is to consolidate the tape drives on an inexpensive loop switch, and then connect the switch to a single director or edge switch E_Port.

Figure 34 illustrates the consolidation of three tape drives through one B_Port connection to a director. Each tape drive has a 30-MBps bandwidth, therefore the sum of the bandwidths of all consolidated servers is slightly less (90 MBps) than the B_Port bandwidth of 1.0625 Gbps. Connecting another FC-AL tape drive to the switch would exceed the B_Port capability and adversely impact director-to-switch link performance. Other devices (such as servers) should not be connected to a switch used for tape drive consolidation.

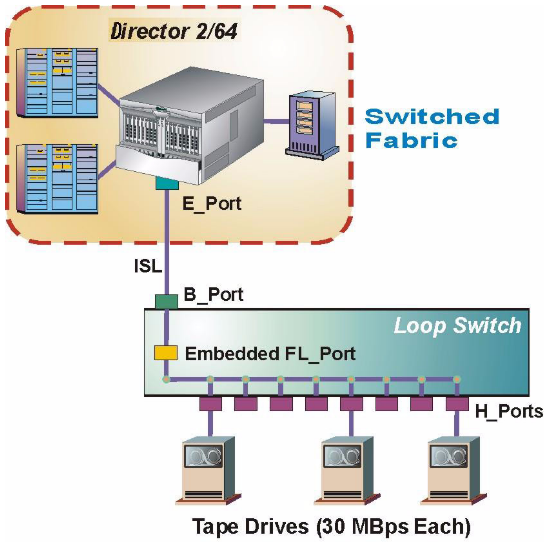

Figure 34: Tape drive consolidation

Planning for Multi-Switch Fabric Support

A Fibre Channel topology that consists of one or more interconnected director or switch elements is called a *fabric*. The product operational software provides the ability to interconnect directors and switches (through E_Port connections) to form a multi-switch fabric. Support of multi-switch fabric operation is a major feature of a director or edge switch.

Consider installation of multiple directors or switches to form a high-availability fabric topology that supports multiple, full-bandwidth data transmission paths between servers and devices. Figure 35 illustrates a simple multi-switch fabric. In the figure, the three fabric elements are Director 2/64s.

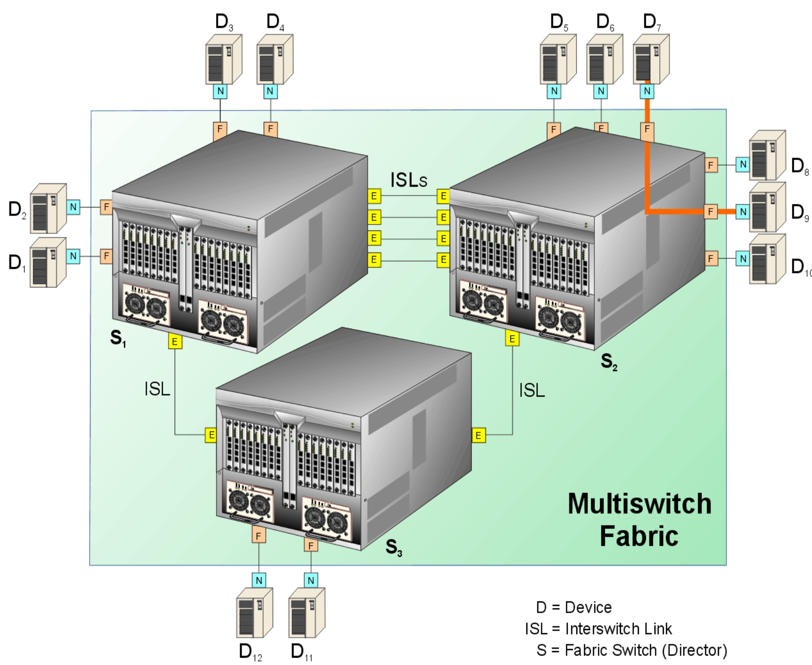

Figure 35: Example multi-switch fabric

Fabric elements cooperate to receive data from the N_Port of an attached device, route the data through the proper director or switch fabric ports (F_Ports), and deliver the data to the N_Port of a destination device. The data transmission path through the fabric is typically determined by the fabric elements and is transparent to the user. Subject to zoning restrictions, devices attached to any of the interconnected directors or switches can communicate with each other through the fabric.

A multi-switch fabric is typically complex and provides the facilities to maintain routing to all device N_Ports attached to the fabric, handle flow control, and satisfy the requirements of the classes of Fibre Channel service that are supported.

Fabric Topology Limits

Operation of multiple directors or switches in a fabric topology is subject to the following topology limits. Consider the impact of these limits when planning the fabric.

- **Fabric Elements** — Each fabric element is defined by a unique domain identification (domain ID) that ranges between 1 and 31. A domain ID of 0 is invalid. Therefore, the theoretical limit of interconnected directors or switches supported in a single fabric is 31.

  For additional information, refer to “Large Fabric Design Implications” on page 109. For the latest supported topology limits, refer to

  http://h18000.www1.hp.com/products/storageworks/san/documentation.html

  or contact your local HP sales representative.

- **Heterogeneous fabric** — Vendor interoperability in the fabric environment is supported; therefore, fabric elements can include directors, edge switches, and open-fabric compliant products supplied by original equipment manufacturers (OEMs). To determine if interoperability is supported for a product, or if communication restrictions apply, refer to the supporting publications for the product or contact your HP representative.

- **Number of ISLs** — The Director 2/64 supports 32 ISLs and the Director 2/140 supports 70 ISLs. Edge Switches 2/12, 2/16, 2/24, and 2/32 support all ports. For redundancy, at least two ISLs should connect any pair of director-class fabric elements. For information, contact your local HP sales representative or refer to the following web site:

  http://h18000.www1.hp.com/products/storageworks/san/documentation.html.

- **Hop count** — The Fibre Channel theoretical limit of ISL connections traversed (hop count) in a single path through the fabric is seven. The maximum hop count supported by a fabric is based on current design rules. For information, refer to

  http://h18000.www1.hp.com/products/storageworks/san/documentation.html

  or contact your local HP sales representative.

Note: The hop count is equal to the number of ISL connections traversed in a single path, not the total number of ISL connections between devices. As shown in Figure 35, the number of ISL connections between Switch S_1 and S_2 is 4, while the number of hops is 1.

Factors to Consider When Implementing a Fabric Topology

Director and switch-based fabrics offer scalable, high-performance, and high-availability connectivity solutions for the enterprise. To enable a multi-switch fabric, all fabric elements must be defined to the *HAFM* application and must be physically cabled to form the requisite ISL connections. In addition, HP recommends that each director or switch in the fabric be assigned a unique preferred domain ID.

When planning to implement a fabric topology, consider the following connectivity and cabling concepts:

- **Physical characteristics and performance objectives** — Most enterprises have unique configurations determined by the characteristics of end devices, fabric elements, cost, and the installation's performance objectives (such as high data transfer rate or high availability). These factors, along with nondisruptive growth and service requirements, must be evaluated when planning an initial fabric. For additional information, refer to "Planning a Fibre Channel Fabric Topology" on page 99.
- **Distance requirements** — The distance between elements in a fabric affects the type of optical port transceiver and cabling required. In addition, variables such as the number of connections, grade of fiber-optic cable, device restrictions, application restrictions, buffer-to-buffer credit limits, and performance requirements can affect distance requirements. Consider the following:
  - If the distance between two fabric elements is less than 250 meters (at 2.125 Gbps), any port type (shortwave or longwave laser) and any fiber-optic cable type (multimode or single-mode) can be used to create an ISL connection. In this case, cost or port availability may be the determining factor.
  - If the distance between two fabric elements exceeds 300 meters (at 2.125 Gbps), only longwave laser ports and single-mode fiber-optic cable can be used to create an ISL.

— Distance limitations can be increased by using multiple fabric elements. Each director or switch retransmits received signals, thus performing a repeater and multiplexer function. Distance limitations can also be increased by using a variety of local area network (LAN), metropolitan area network (MAN), or wide area network (WAN) extension technologies.

Note: Variables such as the number of connections, grade of fiber-optic cable, device restrictions, application restrictions, buffer-to-buffer credit limits, and performance requirements can affect distance requirements.

- **Bandwidth** — ISL connections can be used to increase the total bandwidth available for data transfer between two directors or switches in a fabric. Increasing the number of ISLs between elements increases the corresponding total ISL bandwidth but decreases the number of port connections available to devices. Table 2 illustrates ISL transfer rate versus port availability for a fabric consisting of two Director 2/64s.

Table 2: ISL Transfer Rate Versus Fabric Port Availability (Two-Director Fabric)

Number of ISLs	ISL Data Transfer Rate (at 1.0625 Gbps)	ISL Data Transfer Rate (at 2.125 Gbps)	Available Fabric Ports
1	1.0625 Gbps	2.1250 Gbps	126
2	2.1250 Gbps	4.2500 Gbps	124
3	3.1875 Gbps	6.3750 Gbps	122
4	4.2500 Gbps	8.5000 Gbps	120
5	5.3125 Gbps	10.6250 Gbps	118
6	6.3750 Gbps	12.7500 Gbps	116
7	7.4375 Gbps	14.8750 Gbps	114
8	8.5000 Gbps	17.0000 Gbps	112

- **Load balancing** — Planning consideration must be given to the amount of data traffic expected through the fabric or through a fabric element. Because the fabric automatically determines and uses the least cost (shortest) data transfer path between source and destination ports, some ISL connections may provide insufficient bandwidth while the bandwidth of other connections is unused. To optimize bandwidth use and automatically provide dynamic load balancing across multiple ISLs, consider purchasing and enabling the

OpenTrunking feature key. For information about the feature and managing multiple ISLs, refer to "Open Trunking" on page 167 and "Large Fabric Design Implications" on page 109.

- **Preferred path** — Preferred path is an option that allows a user to configure an ISL data path between multiple fabric elements (directors and switches) by configuring the source and exit ports of the origination fabric element, and the domain ID of the destination fabric element. Each participating director or switch must be configured as part of a desired path. For information about the feature, refer to "Preferred Path" on page 151.

Note: Activating a preferred path can result in receipt of out-of-order frames if the preferred path differs from the current path, if input and output (I/O) is active from the source port, and if congestions is present on the current path.

Fibre Channel frames are routed through fabric paths that implement the minimum possible hop count. For example, in Figure 35, all traffic between devices connected to Director $\mathbf{S_1}$ and Director $\mathbf{S_2}$ communicates directly through ISLs that connect the directors (one hop). No traffic is routed through Director $\mathbf{S_3}$ (two hops). If heavy traffic between the devices is expected, multiple ISL connections should be configured to create multiple minimum-hop paths. With multiple paths, the directors balance the load by assigning traffic from different ports to different minimum-hop paths (ISLs).

When balancing a load across multiple ISLs, a director or switch attempts to avoid assigning multiple ports attached to a device to the same ISL. This minimizes the probability that failure of a single ISL will affect all paths to the device. However, because port assignments are made incrementally as devices log into the fabric and ISLs become available, optimal results are not guaranteed.

Special consideration must also be given to applications with high data transfer rates or devices that participate in frequent or critical data transfer operations. For example, in Figure 35, suppose device $\mathbf{D_7}$ is a server and device $\mathbf{D_9}$ is a storage unit and both devices participate in a critical nightly backup operation. HP recommends that such a connection be routed directly through Director $\mathbf{S_2}$ (rather than the entire fabric) through zoned port connections, WWN-bound port connections, or a preferred path. For additional information, refer to "Device Locality" on page 102.

- **Zoning** — For multi-switch fabrics, zoning is configured on a fabric-wide basis. Changes to the zoning configuration apply to all directors and switches in the fabric. To ensure the zoning configuration is maintained, certain rules are enforced when two or more elements are connected through ISLs to form a fabric, or when two or more fabrics are joined. For additional information, refer to "Configuring Zones" on page 156.

After directors and edge switches are defined and cabled, they automatically join to form a single fabric through a user-transparent process. However, the user should be aware of the following fabric concepts, configuration characteristics, and operational characteristics:

- **Principal Switch selection** — Setting this value determines the principal switch for the multi-switch fabric. Select either **Principal** (highest priority), **Default**, or **Never Principal** (lowest priority) from the **Switch Priority** drop-down list. If all fabric elements are set to **Principal** or **Default**, the director or switch with the highest priority and the lowest WWN becomes the principal switch.

  Following are examples of principal switch selection when fabric elements have these settings.

  — If you have three fabric elements and set all to **Default**, the director or switch with the lowest WWN becomes the principal switch.

  — If you have three fabric elements and set two to **Principal** and one to **Default**, the element with the **Principal** setting that has the lowest WWN becomes the principal switch.

  — If you have three fabric elements and set two to **Default** and one to **Never Principal**, the element with the **Default** setting and the lowest WWN becomes the principal switch.

  Note that at least one director or switch in a multi-switch fabric needs to be set as **Principal** or **Default**. If all the fabric elements are set to **Never Principal**, all ISLs will segment. If all but one element is set to **Never Principal** and the element that was **Principal** goes offline, then all of the other ISLs will segment.

Note: HP recommends configuring the switch priority as **Default**.

In the audit log, note that the Principal setting maps to a number code of 1, Default maps to a number code of 254, and Never Principal maps to a number code of 255. Number codes 2 through 253 are not used.

- **Fabric WWN assignment** — The *Fabric Manager* application identifies fabrics using a fabric WWN. The fabric WWN is the same as the WWN of the fabric's principal switch. If a new principal switch is selected because of a change to the fabric topology, the fabric WWN changes to the WWN of the newly selected principal switch.
- **Domain ID assignment** — Each director or switch in a multi-switch fabric is identified by a unique domain ID that ranges between 1 and 31. A domain ID of 0 is invalid. Numerical domain IDs specified by a user are converted to hexadecimal format and are used in 24-bit Fibre Channel addresses that uniquely identify source and destination ports in a fabric.

  Each fabric element is configured through the *Element Manager* application with a preferred domain ID. When a director or switch powers on and comes online, it requests a domain ID from the fabric's principal switch (indicating its preferred value as part of the request). If the requested domain ID is not allocated to the fabric, the domain ID is assigned to the requesting director or switch. If the requested domain ID is already allocated, an unused domain ID is assigned.

  If two operational fabrics join, they determine if any domain ID conflicts exist between the fabrics. If one or more conflicts exist, the interconnecting ISL E_Ports segment to prevent the fabrics from joining. To prevent this problem, HP recommends that all directors and switches be assigned a unique preferred domain ID. This is particularly important if zoning is implemented through port number (and by default domain ID) rather than WWN.

  When assigning preferred domain IDs in an open fabric with directors and switches supplied by multiple OEMs, be aware of the following:

  — For directors and switches, the firmware adds a base offset of **96** (hexadecimal **60**) to the numerically assigned preferred domain ID. Therefore, if a user assigns a director or switch a numerical preferred domain ID of **1**, the firmware assigns a hexadecimal domain ID of **61**.

  — For non-HP directors and switches, the product firmware may not add a base offset to the numerical preferred domain ID or may add a different hexadecimal base offset (not **20** or **60**).

  As a consequence of this variable base offset and hexadecimal conversion, domain ID conflicts may exist in an open fabric, even if each participating director and switch is assigned a unique numerical domain ID. To determine the method of preferred domain ID assignment for a product, refer to the supporting OEM publications for the product or contact HP.

- **Path selection** — Directors and switches are not manually configured with data transmission paths to each other. Participating fabric elements automatically exchange information to determine the fabric topology and resulting minimum-hop data transfer paths through the fabric. These paths route Fibre Channel frames between devices attached to the fabric and enable operation of the fabric services firmware on each director or switch.

  Paths are determined when the fabric topology is determined and remain static as long as the fabric does not change. If the fabric topology changes (elements are added or removed or ISLs are added or removed), directors and switches detect the change and define new data transfer paths as required. The algorithm that determines data transfer paths is distributive and does not rely on the principal switch to operate. Each director or switch calculates its own optimal paths in relation to other fabric elements.

  Only minimum-hop data transfer paths route frames between devices. If an ISL in a minimum-hop path fails, directors and switches calculate a new least-cost path (which may include more hops) and route Fibre Channel frames over that new path. Conversely, if the failed ISL is restored, directors and switches detect the original minimum-hop path and route Fibre Channel frames over that path.

  When multiple minimum-hop paths (ISLs) between fabric elements are detected, firmware balances the data transfer load and assigns ISLs as follows:

  — The director or switch assigns an equal number of device entry ports (F_Ports) to each E_Port connected to an ISL. For example, if a fabric element has two ISLs and six attached devices, the load from three devices is transferred through each ISL.

  — If a single device has multiple F_Port connections to a director or switch, the switch assigns the data transfer load across multiple ISLs to maximize device availability.

- **Frame delivery order** — When directors or switches calculate a new least-cost data transfer path through a fabric, routing tables immediately implement that path. This may result in Fibre Channel frames being delivered to a destination device out of order, because frames transmitted over the new (shorter) path may arrive ahead of previously transmitted frames that traverse the old (longer) path. This can cause problems because many Fibre Channel devices cannot receive frames in the incorrect order.

Note: Activating a preferred path can result in receipt of out-of-order frames if the preferred path differs from the current path, if input and output (I/O) is active from the source port, and if congestion is present on the current path.

A rerouting delay parameter can be enabled at the *Element Manager* application to ensure the director or switch provides correct frame order delivery. The delay period is equal to the error detect time-out value (E_D_TOV) specified in the *Element Manager* application. Class 2 frames transmitted into the fabric during this delay period are rejected; Class 3 frames are discarded without notification. By default, the rerouting delay parameter is enabled.

Note: To prevent E_Port segmentation, the same E_D_TOV and resource allocation time-out value (R_A_TOV) must be specified for each fabric element.

- **E_Port segmentation** — When an ISL activates, the two fabric elements exchange operating parameters to determine if they are compatible and can join to form a single fabric. If the elements are incompatible, the connecting E_Port at each director or switch segments to prevent the creation of a single fabric. A segmented link transmits only Class F traffic; the link does not transmit Class 2 or Class 3 traffic.

  The following conditions cause E_Ports to segment:

  - **Incompatible operating parameters** — Either the R_A_TOV or E_D_TOV is inconsistent between the two fabric elements.
  - **Duplicate domain IDs** — One or more domain ID conflicts are detected.
  - **Incompatible zoning configurations** — Zoning configurations for the two fabric elements are not compatible. For an explanation, refer to "Configuring Zones" on page 156.
  - **Build fabric protocol error** — A protocol error is detected during the process of forming the fabric.
  - **No principal switch** — No director or switch in the fabric is capable of becoming the principal switch.

— **No response from attached switch** — After a fabric is created, each element in the fabric periodically verifies operation of all attached switches and directors. An ISL segments if a director or switch does not respond to a verification request.

— **ELP retransmission failure timeout** — A director or switch that exhibits a hardware failure or connectivity problem cannot transmit or receive Class F frames. The director or switch did not receive a response to multiple exchange link parameters (ELP) frames, did not receive a fabric login (FLOGI) frame, and cannot join an operational fabric.

- **Fabric services and state change notifications** — In a multi-switch fabric, services provided by each director or switch (such as name service, registered state change notifications [RSCNs], and zoning) are provided on a fabric-wide basis. For example, if a fabric-attached device queries a director or switch name server to locate all devices that support a specified protocol, the reply includes all fabric devices that support the protocol that are in the same zone as the requesting device, not just devices attached to the director or switch.

  RSCNs are transmitted to all registered device N_Ports attached to the fabric if either of the following occur:

  — A fabric-wide event occurs, such as a director or switch logging in to the fabric, a director or switch logging out of the fabric, or a reconfiguration because of a director, switch, or ISL failure.

  — A zoning configuration changes.

- **Zoning configurations for joined fabrics** — In a multi-switch fabric, zoning is configured on a fabric-wide basis, and any change to the active zone set is applied to all directors and switches. To ensure zoning is consistent across a fabric, the following rules are enforced when two fabrics (zoned or unzoned) join through an ISL.

  — **Fabric A unzoned and Fabric B unzoned** — The fabrics join successfully, and the resulting fabric remains unzoned.

  — **Fabric A zoned and Fabric B unzoned** — The fabrics join successfully, and fabric B automatically inherits the zoning configuration from fabric A.

  — **Fabric A unzoned and Fabric B zoned** — The fabrics join successfully, and fabric A automatically inherits the zoning configuration from fabric B.

— **Fabric A zoned and Fabric B zoned** — The fabrics join successfully only if the zone sets can be merged. If the fabrics cannot join, the connecting E_Ports segment and the fabrics remain independent.

Zone sets for two directors or switches are compatible (the fabrics can join) only if the zone names for each fabric element are unique. The zone names for two fabric elements can be the same only if the zone member WWNs are identical for each duplicated zone name.

Fabric Topologies

Several topologies exist from which to build a Fibre Channel fabric infrastructure. This section describes the most effective fabric topologies and provides guidance on when to deploy each topology. The topologies are effective for a wide variety of applications, are extensively tested by HP, and are deployed in several customer environments.

Fabric topologies described in this section include:

- Cascaded Fabric
- Ring Fabric
- Mesh Fabric
- Core-to-Edge Fabric
- Fabric Island

Cascaded Fabric

A cascaded fabric consists of a linear string of directors or switches connected by one or more ISLs. Each fabric element is connected to the next fabric element in line. The end-point fabric elements are not connected to each other. Figure 36 illustrates a cascaded fabric topology.

Cascaded fabrics are typically inexpensive, easy to deploy, and provide a simple solution to add additional fabric devices. However, this fabric design has low reliability because each director, switch, or ISL is a single point of failure. In addition, the design has limited scalability because the maximum hop count can be quickly exceeded when fabric elements are added.

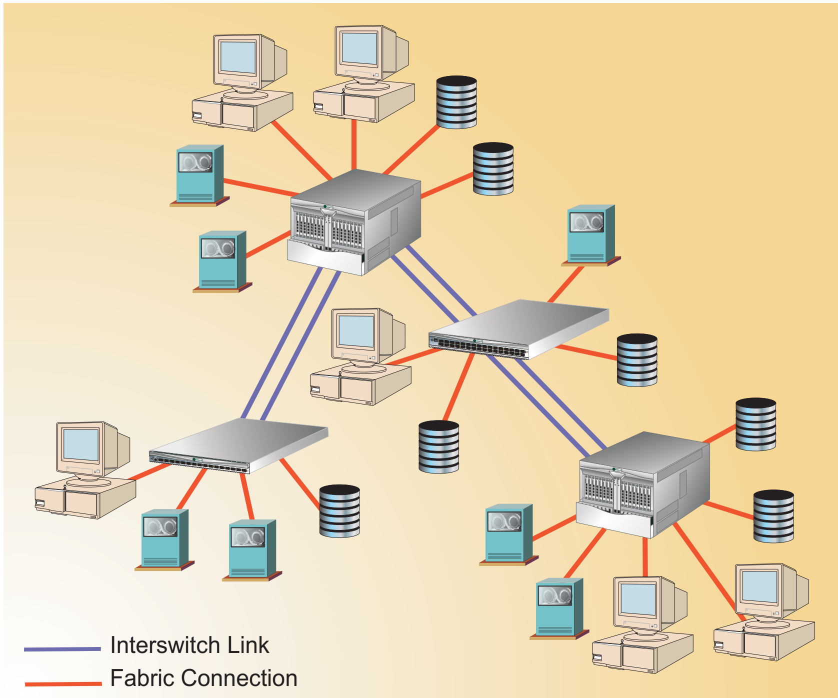

Figure 36: Cascaded fabric

One design variation is to use more than one ISL between fabric elements. This eliminates ISLs as a single point of failure and greatly increases the fabric design reliability.

Cascaded fabrics are well suited for applications where data access is local but not for applications that require any-to-any connectivity. Device locality implies that groups of servers and the storage they access are connected through the same fabric element and that ISLs are used primarily for fabric management traffic (Class F traffic) or low-bandwidth SAN applications. For additional information, refer to "Device Locality" on page 102.

Ring Fabric

A ring fabric consists of a continuous string of directors or switches connected by one or more ISLs. Each fabric element is connected to the next fabric element (like a cascaded fabric, but with the end-point fabric elements connected). Figure 37 illustrates a ring fabric topology.

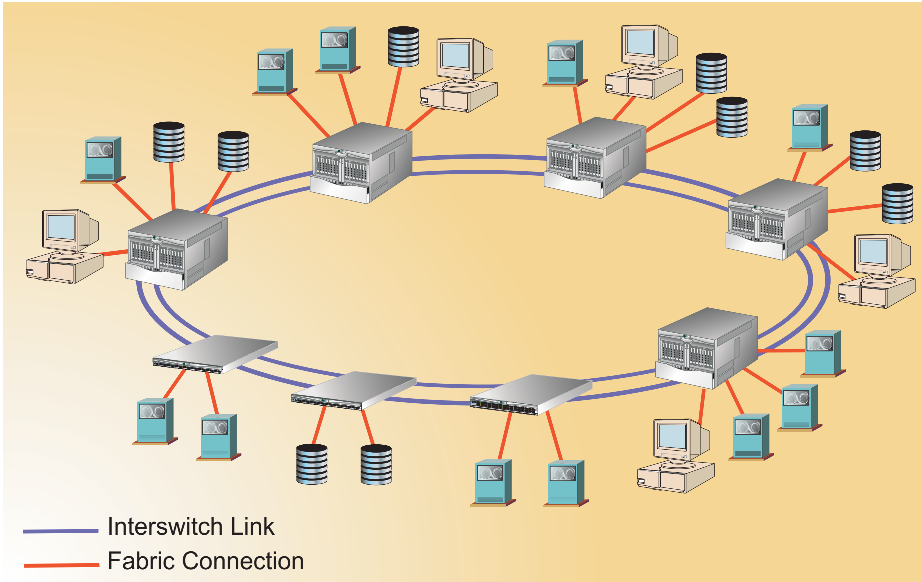

Figure 37: Ring fabric

Ring fabrics are generally more expensive than cascaded fabrics, are also easy to deploy, provide a simple solution to add additional fabric devices, and can solve hop-count problems inherent to cascaded fabrics. Ring fabrics also have increased reliability because traffic can route around a single ISL, director, or switch failure (subject to hop count limitations).

Like cascaded fabrics, ring fabrics are well suited for applications where data access is local, but not for applications that require any-to-any connectivity. In addition, ring fabrics are useful when connecting SANs over a MAN or WAN. These networks typically use a ring topology.

Mesh Fabric

There are two types of mesh fabrics: full mesh and modified (or partial) mesh. In a full-mesh topology, every director or switch is directly connected to all directors and switches in the fabric. The maximum hop count between fabric-attached devices is one hop. Figure 38 illustrates a full-mesh fabric topology.

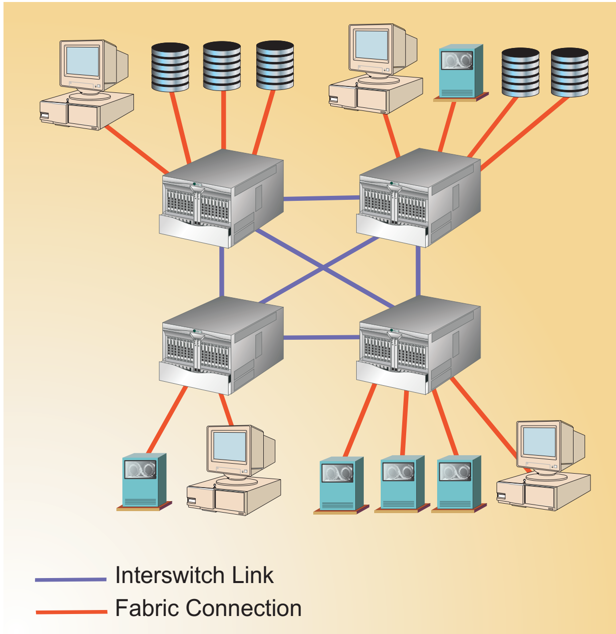

Figure 38: Full mesh fabric

Full-mesh fabrics provide increased resiliency over cascaded or ring fabrics and are well suited for applications that require any-to-any connectivity. If a single ISL fails, traffic is automatically routed through an alternate path.

Mesh fabrics also form effective backbones to which other SAN islands can be connected. Traffic patterns through the fabric should be evenly distributed, and overall bandwidth consumption is low.

When using low port-count fabric elements, mesh fabrics are best used when the fabric is not expected to grow beyond four or five switches. The cost of ISLs becomes prohibitive for larger mesh fabrics. In addition, full-mesh fabrics do not scale easily because the addition of a switch requires that at least one additional ISL be added from every existing switch in the fabric. If less than four fabric elements are used in a full-mesh fabric:

- A two-switch full mesh fabric is identical to a two-switch cascaded fabric.
- A three-switch full mesh fabric is identical to a three-switch ring fabric.

A modified or partial-mesh fabric is similar to a full-mesh fabric, but each switch does not have to be directly connected to every other switch in the fabric. The fabric is still resilient to failure but does not carry a cost premium for unused or redundant ISLs. In addition, partial-mesh fabrics scale easier than full-mesh fabrics.

Partial-mesh fabrics are useful when designing a SAN backbone for which traffic patterns between SAN islands connected to the backbone are well known. If heavy traffic is expected between a pair of switches, the switches are connected through at least one ISL; if minimal traffic is expected, the switches are not connected.

In general, mesh fabrics can be difficult to scale without downtime. The addition of directors or switches usually involves disconnecting fabric devices and may involve disconnecting in-place ISLs. As a result, full or partial-mesh fabrics are recommended for networks that change infrequently or have well-established traffic patterns.

Core-to-Edge Fabric

A core-to-edge fabric consists of one or more Fibre Channel directors or switches acting as core elements that are dedicated to connecting other directors and switches (edge elements) in the fabric. Core directors act as high-bandwidth routers with connectivity to edge fabric elements. Figure 39 illustrates the core-to-edge fabric topology with two core directors and fourteen edge directors and switches (2-by-14 topology).

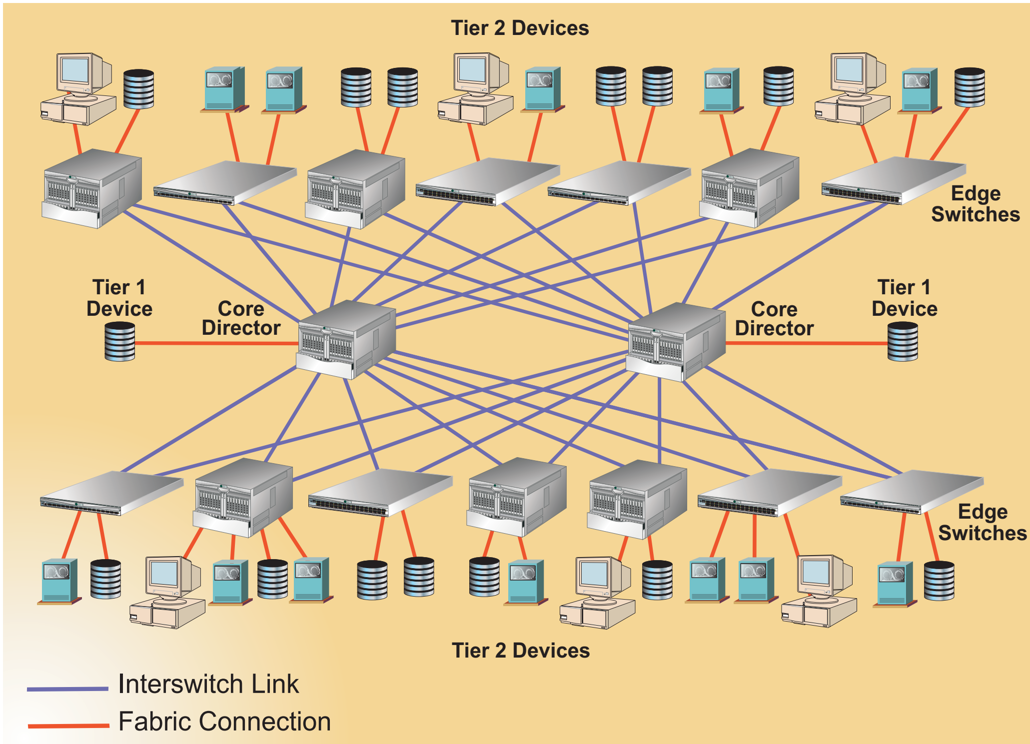

Figure 39: 2-by-14 Core-to-Edge fabric

Subject to large fabric design constraints, core-to-edge fabrics are easy to scale through the addition of core elements. Figure 40 illustrates a core-to-edge fabric topology with four core directors and twelve edge directors and switches (4-by-12 topology).

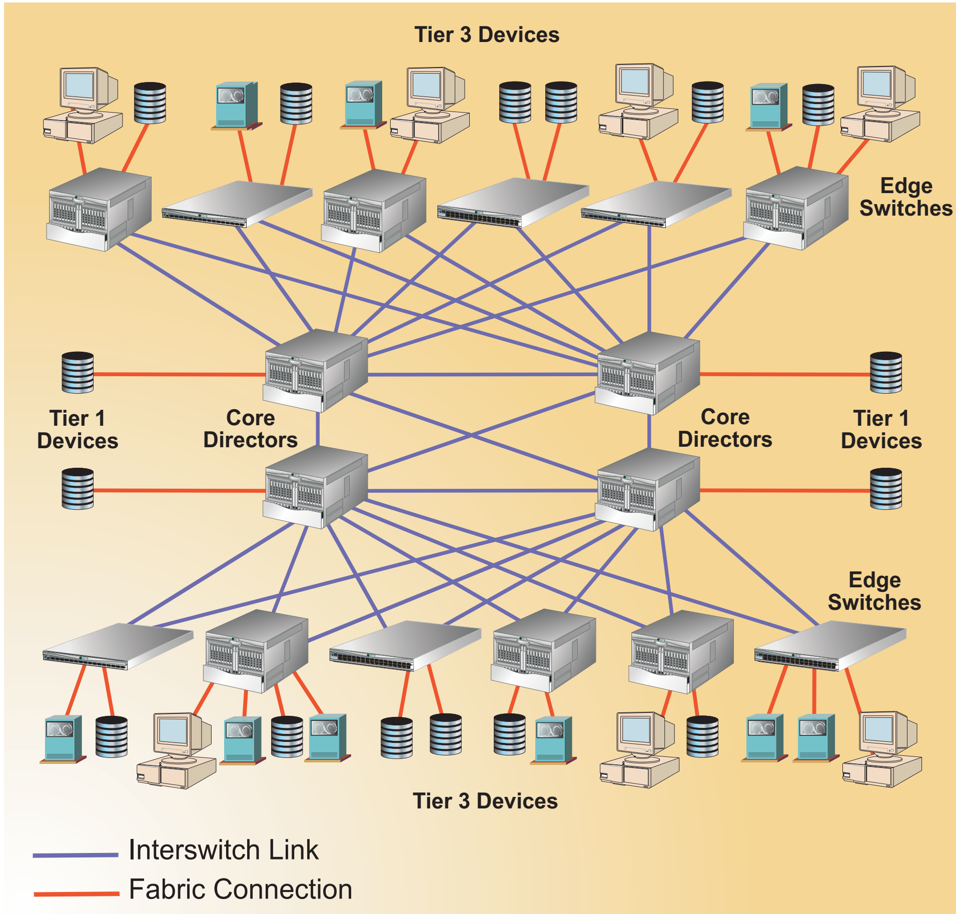

Figure 40: 4-by-12 Core-to-Edge fabric

A core-to-edge topology offers any-to-any device connectivity, and evenly distributes traffic bandwidth throughout the fabric. The topology provides the most flexible architecture to address fabric performance, traffic locality, data integrity, connectivity, and scalability requirements.

The simplest core-to edge fabric has two or more core switching elements that may or may not be connected (simple or complex). In a simple core topology, as shown in Figure 39, core switches are not connected. In a complex core topology, as shown in Figure 40, core switches are connected. The figure also illustrates a topology where the core is a full-mesh fabric.

Each edge switch connects (through at least one ISL) to each core switch but not to other edge switches. There are typically more device connections to an edge switch than ISL connections; therefore, edge switches act as consolidation points for servers and storage devices. The ratio of ISLs to device connections for each switch is a function of device performance. For additional information, refer to "FCP and FICON in a Single Fabric" on page 110.

Fibre channel devices (servers and storage devices) connect to core or edge fabric elements in tiers. These tiers are defined as follows:

- **Tier 1** — A Tier 1 device connects directly to a core director or switch. Tier 1 devices are typically high-use or high-I/O devices that consume substantial bandwidth and should not be connected through an ISL. In addition, IBM fiber connection (FICON) devices cannot communicate through E_Ports (ISLs) and must use Tier 1 connectivity. For additional information, refer to "FCP and FICON in a Single Fabric" on page 110.
- **Tier 2** — A Tier 2 device connects to an edge switch and Fibre Channel traffic from the device must traverse only one ISL (hop) to reach a device attached to a core director or switch.
- **Tier 3** — A Tier 3 device connects to an edge switch and Fibre Channel traffic from the device can traverse two ISLs (hops) to reach a device attached to a core director or switch.

Fabric Island

A fabric island topology connects several geographically diverse Fibre Channel fabrics. These fabrics may also comprise different topologies (cascaded, ring, mesh, or core-to-edge), but may require connectivity for shared data access, resource consolidation, data backup, remote mirroring, or disaster recovery.

When connecting multiple fabrics, data traffic patterns and fabric performance requirements must be well known. Fabric island connectivity must adhere to topology limits, including maximum number of fabric elements and ISL hop count. It is also essential to maintain data locality within fabric islands as much as possible, and to closely monitor bandwidth usage between the fabric islands.

Planning a Fibre Channel Fabric Topology

To be effective, the fabric topology design must:

- Solve the customer's business problem and provide the required level of performance.
- Meet the customer's requirements for high availability.
- Be scalable to meet future requirements.

Fabric Performance

During the design phase of a Fibre Channel fabric, performance requirements of the fabric and of component directors, switches, and devices must be identified and incorporated. An effective fabric design can accommodate changes to performance requirements, and incorporate additional directors, switches, devices, ISLs, and higher speed links with minimal impact to fabric operation. Performance factors that affect fabric design include:

- Application input/output (I/O) requirements, both in Gbps and I/Os per second (IOPS).
- Storage port fan-out.
- Hardware limits, including the maximum directors and switches per fabric, maximum number of ISLs per director or switch, and maximum hops between devices. For additional information, refer to "Fabric Topology Limits" on page 81.
- Software limits, including the maximum number of fabric elements managed by the *HAFM* application, and the maximum number of zones and zone members. For additional information, refer to "Product Software" on page 47 and "Configuring Zones" on page 156.

I/O Requirements

HP directors and switches are designed with non-blocking architecture; therefore, any two switch ports can communicate at the full Fibre Channel bandwidth of 2.125 Gbps without impact to other switch ports. Because most SAN-attached devices are not capable of generating I/O traffic at the full bandwidth, there is little potential for congestion between two devices attached through a single director or switch.

However, when multiple directors or switches are connected through a fabric ISL that multiplexes traffic from several devices, significant potential for congestion arises. To minimize congestion, factors such as application I/O profiles, ISL oversubscription, and device locality must be included in the fabric design.

Application I/O Profiles

Understanding application I/O characteristics is essential to SAN, fabric, and ISL design. Factors that may affect application I/O include:

- **Read/write mixture** — Although application I/O is typically a mixture of read and write operations, some applications are very biased. For example, video server applications are almost 100% read intensive, while real-time video editing applications are mostly write intensive. Read operations typically take less time than write operations; therefore, storage devices for a read-intensive application usually wait for data transfer. As a consequence, read-intensive applications typically require high bandwidth to the device.
- **Type of data access** — When an application requires data, access to that data is random or sequential. For example, e-mail server activity is random access, while seismic data processing for the oil and gas industry is sequential access. Sequential data access typically takes less time than random data access; therefore, sequential-access applications usually wait for data transfer. As a consequence, sequential-access applications typically require high bandwidth to the device.
- **I/O block size** — The third characteristic of application I/O is data block size, which typically ranges from two kilobytes (KB) to over one megabyte (MB). Applications that generate large blocks of data require high bandwidth to the device.

Prior to fabric design, application I/O profiles should be estimated or established that classify the application bandwidth requirements. Bandwidth consumption is classified as *light*, *medium*, or *heavy*. These classifications must be considered when planning ISL and device connectivity. For information about application I/O (in Gbps) and fabric performance problems due to ISL connectivity, refer to "ISL Oversubscription" on page 101. For information about application I/O (in IOPS) and fabric performance problems due to port contention, refer to "Device Fan-Out Ratio" on page 103.

ISL Oversubscription

ISL oversubscription (or congestion) occurs when multiplexed traffic from several devices is transmitted across a single ISL. When an ISL is oversubscribed, fabric elements use fairness algorithms to interleave data frames from multiple devices, thus giving fractional bandwidth to the affected devices. Although all devices are serviced, ISL and fabric performance is reduced.

Figure 41 illustrates ISL oversubscription. Two NT servers, each with maximum I/O of 100 MBps, are contending for the bandwidth of a single ISL operating at 1.0625 Gbps. In addition to data, the ISL must also transmit Class F traffic internal to the fabric. When operating at peak load, each NT server receives less than half the available ISL bandwidth.

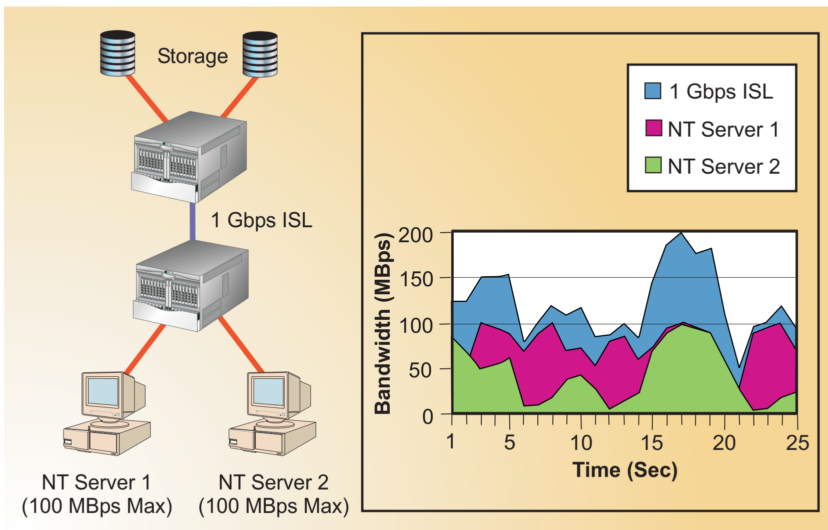

Figure 41: ISL oversubscription

Depending on fabric performance requirements and cost, there are several options to solve ISL oversubscription problems, including:

- **Employ device locality** — NT server 1 and its associated storage device can be connected through one director. NT server 2 and its associated storage device can be connected through the other director. As a result, minimal traffic flows across the ISL between directors and the congestion problem is mitigated. For additional information, refer to "Device Locality" on page 102.

- **Install an additional ISL** — A second ISL can be installed to balance the traffic load between fabric elements. Two ISLs are sufficient to support the bandwidth of both NT servers operating at peak load.
- **Upgrade the existing ISL** — Fabric element software, firmware, and hardware can be upgraded to support a 2.125 Gbps bandwidth traffic load between fabric elements. A 2.125 Gbps ISL is sufficient to support the bandwidth of both NT servers operating at peak load.
- **Deliberately employ ISL oversubscription** — Real-world SANs are expected to function well, even with oversubscribed ISLs. Device I/O is typically bursty; few devices operate at peak load for a significant length of time, and device loads seldom peak simultaneously. As a result, ISL bandwidth is usually not fully allocated, even for an oversubscribed link. An enterprise can realize significant cost savings by deliberately designing a SAN with oversubscribed ISLs that provide connectivity for noncritical applications.

Device Locality

Devices that communicate with each other through the same director or switch have *high locality*. Devices that must communicate with each other through one or more ISLs have *low locality*. Part (A) of Figure 42 illustrates high device locality with little ISL traffic. Part (B) of Figure 42 illustrates low device locality.

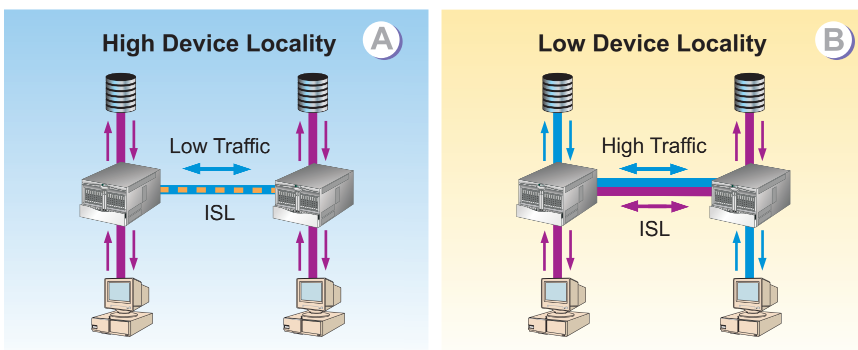

Figure 42: Device locality

Although it is possible to design a SAN that delivers sufficient ISL bandwidth in a zero-locality environment, it is preferable to design local, one-to-one connectivity for heavy-bandwidth applications such as video server, seismic data processing, or medical 3D imaging.

When designing a core-to-edge fabric, servers and storage devices that support such bandwidth-intensive applications should be attached to core directors as Tier 1 devices. As a best practices policy (assuming 1.0625 Gbps ISLs), devices that generate a sustained output of 35 MBps or higher are candidates for Tier 1 connectivity. IBM FICON devices also must use Tier 1 connectivity. For additional information, refer to "FCP and FICON in a Single Fabric" on page 110.

Device Fan-Out Ratio

The output of most host devices is bursty in nature; most devices do not sustain full-bandwidth output, and it is uncommon for the output of multiple devices to peak simultaneously. These variations are why multiple hosts can be serviced by a single storage port. This device sharing leads to the concept of *fan-out ratio*.

Device fan-out ratio is defined as the storage or array port IOPS divided by the attached host IOPS, rounded down to the nearest whole number. A more simplistic definition for device fan out is the ratio of host ports to a single storage port. Fan-out ratios are typically device dependent. In general, the maximum device fan-out ratio supported is 12 to 1. Figure 43 illustrates a fan-out ratio of 10 to 1.

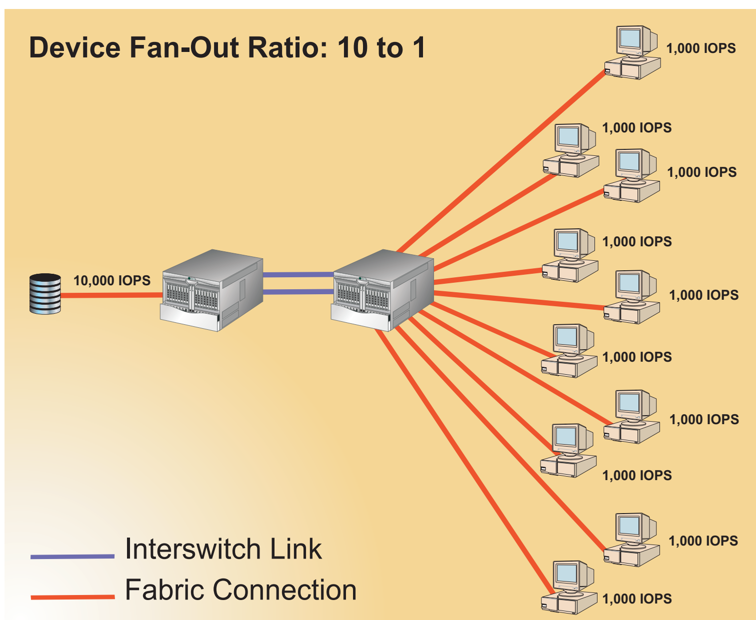

Figure 43: Device fan-out ratio

Performance Tuning

When designing or tuning a fabric for performance, it is critical to understand application I/O characteristics so that:

- Device output in Gbps does not oversubscribe ISLs, leading to fabric congestion.
- Device output in IOPS does not result in a connectivity scheme that exceeds fan-out ratios, leading to port congestion.

Figure 44 illustrates performance tuning for a simple fabric using appropriate ISL connectivity, device locality, and fan-out regions for device connectivity.

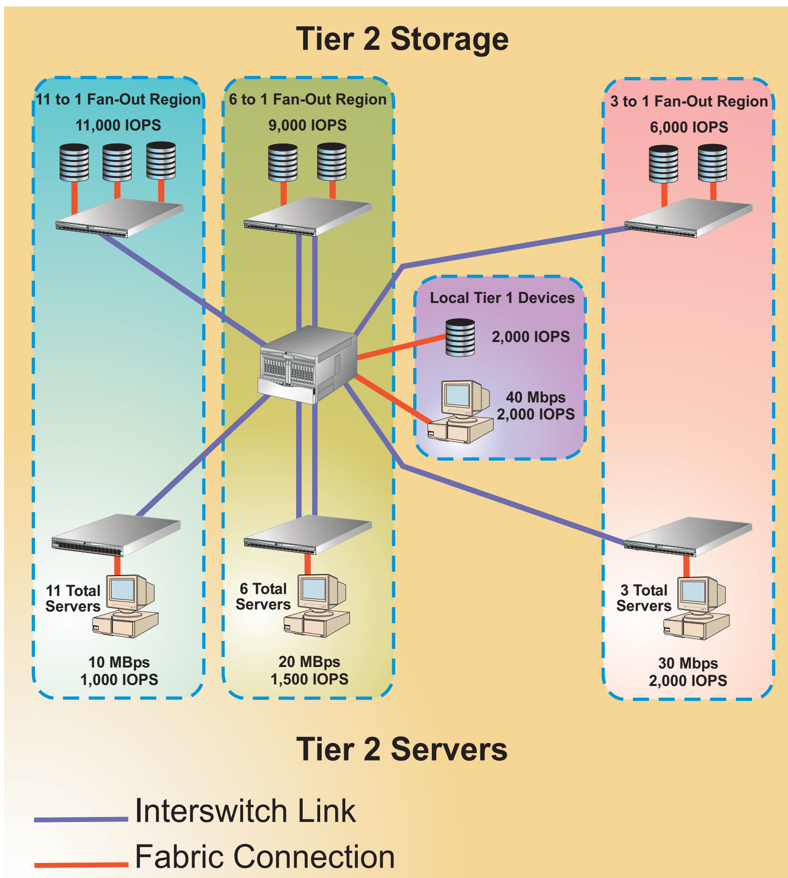

Figure 44: Fabric performance tuning

- **Local Tier 1 devices** — A video server application with I/O capabilities of 40 MBps and 2,000 IOPS must be connected to the fabric. Because the application is critical and high bandwidth (in excess of 35 MBps), the server and associated storage are directly attached to the core director as Tier 1 devices. No ISLs are used for server-to-storage connectivity.
- **11 to 1 fan-out region** — Eleven NT servers with I/O capabilities of 10 MBps and 1,000 IOPS are fabric-attached through a 32-port edge switch. The primary applications are e-mail and online transaction processing (OLTP). Because bandwidth use is light and noncritical, the servers are connected to the core director with a single ISL that is intentionally oversubscribed (1.1 Gbps plus Class F traffic). The servers are connected to storage devices with I/O capabilities of 11,000 IOPS.
- **6 to 1 fan-out region** — Six servers with I/O capabilities of 20 MBps and 1,500 IOPS are fabric-attached through a 16-port edge switch. Bandwidth use is light to medium but critical, so the servers are connected to the core director with two ISLs (0.6 Gbps each plus Class F traffic). The servers are connected to storage devices with I/O capabilities of 9,000 IOPS.
- **3 to 1 fan-out region** — Three servers with I/O capabilities of 30 MBps and 2,000 IOPS are fabric-attached through a 16-port edge switch. Bandwidth use is medium but non-critical, so the servers are connected to the core director with one ISL (0.9 Gbps plus Class F traffic). The servers are connected to storage devices with I/O capabilities of 6,000 IOPS.

Fabric Availability

Many fabric-attached devices require highly available connectivity to support applications such as disk mirroring, server clustering, or business continuance operations. High availability is accomplished by deploying a *resilient fabric topology* or *redundant fabrics*.

A fabric topology that provides at least two internal routes between fabric elements is considered resilient. A single director, switch, or ISL failure does not affect the remaining elements, and the overall fabric remains operational. However, unforeseen events such as human error, software failure, or disaster can cause the failure of a single resilient fabric. Using redundant fabrics (with resiliency) mitigates these effects and significantly increases fabric availability.

Fibre Channel fabrics are classified by four levels of resiliency and redundancy. From least available to most available, the classification levels are:

- **Nonresilient single fabric** — Directors and switches are connected to form a single fabric that contains at least one single point of failure (fabric element or ISL). Such a failure causes the fabric to fail and segment into two or more smaller fabrics. A cascaded fabric topology (Figure 36) illustrates this design.
- **Resilient single fabric** — Directors and switches are connected to form a single fabric, but no single point of failure can cause the fabric to fail and segment into two or more smaller fabrics. A ring fabric topology (Figure 37) illustrates this design.
- **Nonresilient dual fabric** — Half the directors and switches are connected to form one fabric, and the remaining half are connected to form an identical but separate fabric. Servers and storage devices are connected to both fabrics. Each fabric contains at least one single point of failure (fabric element or ISL). All applications remain available, even if an entire fabric fails.
- **Resilient dual fabric** — Half the directors and switches are connected to form one fabric, and the remaining half are connected to form an identical but separate fabric. Servers and storage devices are connected to both fabrics. No single point of failure can cause either fabric to fail and segment. All applications remain available, even if an entire fabric fails and elements in the second fabric fail.

  A dual-fabric resilient topology is generally the best design to meet high-availability requirements. Another benefit of the design is the ability to proactively take one fabric offline for maintenance without disrupting SAN operations.

Redundant Fabrics

If high availability is important enough to require dual-connected servers and storage, a dual-fabric solution is generally preferable to a dual-connected single fabric. Dual fabrics maintain simplicity and reduce (by 50%) the size of fabric routing tables, name server tables, updates, and Class F management traffic. In addition, smaller fabrics are easier to analyze for performance, to fault isolate, and to maintain.

Figure 45 illustrates simple redundant fabrics. Fabric “A” and fabric “B” are symmetrical, each containing one core director and four edge switches. All servers and storage devices are connected to both fabrics.

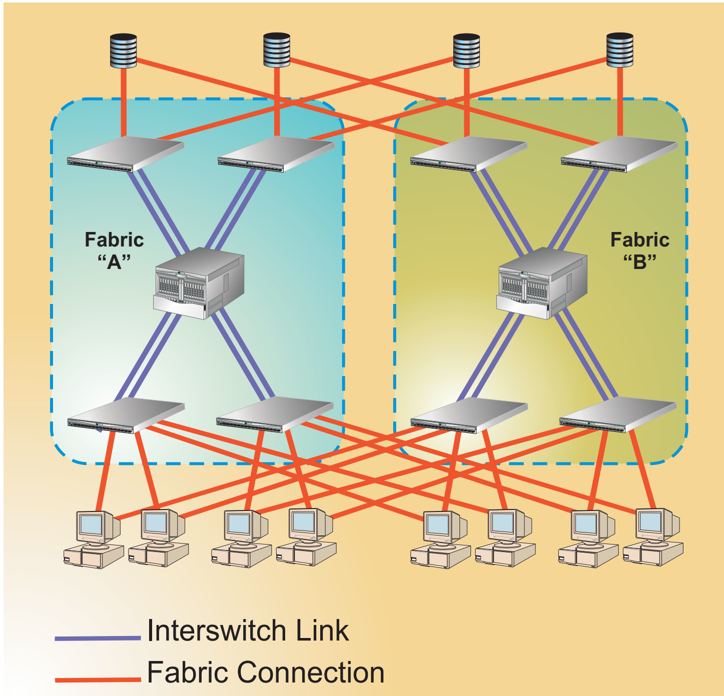

Figure 45: Redundant fabrics

Some dual-attached devices support active-active paths, while others support only active-passive paths. Active-active devices use either output path equally, and thus use both fabrics and double the device bandwidth. Active-passive devices use the passive path only when the active path fails.

When deploying redundant fabrics, it is not required that the fabrics be symmetrical. As an example, single-attached devices, such as tape drives and noncritical servers and storage, can be logically grouped and attached to one of the fabrics.

Fabric Scalability

A scalable fabric allows for nondisruptive addition of fabric elements (directors and switches) or ISLs to increase the size or performance of the fabric. Scalability also relates to investment protection. If a core edge switch is replaced with a newer or higher port count switch, it is often valuable to use the existing switch elsewhere in the fabric (at the edge).

Obtaining Professional Services

Planning and implementing a multi-switch fabric can be a complex and difficult task. HP recommends that you obtain planning assistance from our professional services organization before implementing a fabric topology.

Fabric Topology Design Considerations

This section discusses additional fabric topology design considerations, including:

- Large Fabric Design Implications
- FCP and FICON in a Single Fabric
- Multiple Data Transmission Speeds in a Single Fabric
- Fibre Channel Distance Extension

Large Fabric Design Implications

Businesses are experiencing an unprecedented growth of information and the requirement to maintain that information online. To meet these requirements, Fibre Channel SANs provide the infrastructure to connect thousands of servers to hundreds of storage devices. To provide enterprise-class SAN performance and scalability, large fabric designs are required.

When multiple directors or switches are connected, ISL (E_Port) communication must be established between fabric elements and the fabric must be initialized. During fabric initialization, the fabric elements:

- Establish the operating mode for connected E_Port pairs and exchange link parameters (E_Port names, timeout values, class-specific information, and flow control parameters).
- Exchange fabric parameters, select a principal switch, and assign domain IDs to all switches.
- Employ a routing protocol to establish the shortest path through the fabric and program route tables for each fabric element.
- Exchange the active zone set to ensure uniform zoning is enforced between all fabric elements.

However, fabric initialization is not a serial process. The process executes concurrently across all ISLs in the fabric, causing a massive flood of Class F traffic that must be processed to the embedded port of each fabric element within a specified (fabric-wide) error detect time-out value (E_D_TOV). If the fabric consists of a large number of elements (and therefore ISLs), Class F traffic may not be processed within the E_D_TOV, resulting in error recovery operations, time-outs, segmented links, or fabric failure.

Because of these problems, a fabric with a high ISL count is more difficult to build. Note that the fabric problem is not directly related to the large number of fabric elements but to the large number of ISLs associated with the elements. Fabric build concerns currently limit the combined number of directors and switches to about 24.

FCP and FICON in a Single Fabric

Fibre Channel Layer 4 (FC-4) describes the interface between Fibre Channel and various upper-level protocols. FCP and FICON are the major FC-4 protocols. FCP is the Fibre Channel protocol that supports the small computer system interface (SCSI) upper-level transport protocol. FICON is the IBM successor to the enterprise systems connection (ESCON) protocol and adds increased reliability and integrity to that provided by the FCP protocol.

Because FCP and FICON are both FC-4 protocols, routing of Fibre Channel frames is not affected when the protocols are mixed in a single fabric environment. However, management differences in the protocols arise when a user changes director or edge switch parameters through zoning or connectivity control. In particular:

- FCP communication parameters are port number and name-centric, discovery-oriented, and assigned by the fabric, and they use the Fibre Channel name server to control device communication.
- FICON communication parameters are logical port address-centric, definition oriented, and assigned by the attached host and they use host assignment to control device communication.

In addition to OEM limitations not discussed in this publication, the considerations that need to be evaluated when intermixing FCP and FICON protocols are:

- Director or Switch Management
- Port Numbering Versus Port Addressing
- Management Limitations
- Features that Impact Protocol Intermixing
- Protocol Intermixing Best Practices

Director or Switch Management

When intermixing FCP and FICON protocols, it must be determined if the director or switch is to be managed through Open Systems management style or FICON management style. This setting affects only the management style used to manage the director or switch; it does not affect Fibre Channel port operation. FCP devices can communicate with each other when the attached fabric element is set to FICON management style, and FICON devices can communicate with each other when the attached fabric element is set to Open Systems management style.

When a director or switch is set to Open Systems management style, FCP connectivity is defined within a Fibre Channel fabric using WWNs of devices that are allowed to form connections. When connecting to the fabric, an FCP device queries the name server for a list of devices for which connectivity is allowed. This connectivity is hardware enforced through a name server zoning feature that partitions attached devices into restricted-access zones.

When a director or switch is set to FICON management style, host-to-storage FICON connectivity and channel paths are defined by a host-based hardware configuration definition (HCD) program, a director or switch-resident management server called the control unit port (CUP), and a user-configured (and director or switch-resident) prohibit dynamic connectivity mask (PDCM) array associated with each logical port address. FICON devices do not query the name server for accessible devices because connectivity is defined at the host. This connectivity is hardware enforced in the routing tables of each port.

Note: The Edge Switch 2/12 and Edge Switch 2/24 do not support operation using FICON management style or the transmission of FICON frames.

PDCM connectivity control is configured and managed at the director or switch level using the **Configure Allow/Prohibit Matrix - Active** dialog box (Figure 57). For additional information, refer to "PDCM Arrays" on page 149.

Note: When configuring a PDCM array that prohibits E_Port connectivity, mistakes can render ISLs unusable and cause complex routing problems. These problems can be difficult to fault isolate and sometimes manifest incorrectly as end-device issues.

Port Numbering Versus Port Addressing

Consideration must be given to the implications of port numbering for the FCP protocol versus logical port addressing for the FICON protocol. FCP configuration attributes are implemented through zoning. Zones are configured through the associated *Element Manager* application by authorizing or restricting access to name server information associated with device N_Ports that attach to director or switch F_Ports.

Zones are configured by:

- The 8-byte (64-digit) WWN assigned to the host bus adapter (HBA) or Fibre Channel interface installed in a device connected to the director or switch.
- The domain identification (ID) and physical port number of the director or edge switch port to which a device is attached.

FICON configuration attributes are implemented through logical port addressing. This concept is consistent with the address-centric nature of FICON and allows ports to be swapped for maintenance operations without regenerating a host configuration.

Logical port addresses are derived by converting the port number from numerical to hexadecimal format and adding a hexadecimal 4 to the result. Figure 46 illustrates port numbering and logical port addressing for the Director 2/64. The figure shows:

- Universal port module (UPM) card numbers at the top (numerical **0** through **15**).
- Numerical physical port numbers in blue (**00** through **63**).
- Hexadecimal physical port numbers in red (**00** through **3F**).
- Logical port addresses in bold (hexadecimal **04** through **43**).

UPM Cards								CTP2 - 1 Card	CTP2 - 0 Card	UPM Cards							
15	**14**	**13**	**12**	**11**	**10**	**9**	**8**			**7**	**6**	**5**	**4**	**3**	**2**	**1**	**0**
63 3F **43**	59 3B **3F**	55 37 **3B**	51 33 **37**	47 2F **33**	43 2B **2F**	39 27 **2B**	35 23 **27**			31 1F **23**	27 1B **1F**	23 17 **1B**	19 13 **17**	15 0F **13**	11 0B **0F**	07 07 **0B**	03 03 **07**
62 3E **42**	58 3A **3E**	54 36 **3A**	50 32 **36**	46 2E **32**	42 2A **2E**	38 26 **2A**	34 22 **26**			30 1E **22**	26 1A **1E**	22 16 **1A**	18 12 **16**	14 0E **12**	10 0A **0E**	06 06 **0A**	02 02 **06**
61 3D **41**	57 39 **3D**	53 35 **39**	49 31 **35**	45 2D **31**	41 29 **2D**	37 25 **29**	33 21 **25**			29 1D **21**	25 19 **1D**	21 15 **19**	17 11 **15**	13 0D **11**	09 09 **0D**	05 05 **09**	01 01 **05**
60 3C **40**	56 38 **3C**	52 34 **38**	48 30 **34**	44 2C **30**	40 28 **2C**	36 24 **28**	32 20 **24**			28 1C **20**	24 18 **1C**	20 14 **18**	16 10 **14**	12 0C **10**	08 08 **0C**	04 04 **08**	00 00 **04**

Figure 46: Director 2/64 port numbers and logical port addresses

Although Figure 46 depicts a UPM card map only for the Director 2/64, physical port numbers and logical port addresses can be extrapolated for the Director 2/140 (140 ports), Edge Switch 2/12 (12 ports), Edge Switch 2/16 (16 ports), Edge Switch 2/24 (24 ports), and Edge Switch 2/32 (32 ports).

Management Limitations

The following considerations must be given to the limitations and interactions of director or switch management when using Open Systems management style (FCP) or FICON management style:

- FICON port-to-port connectivity is hardware enforced, while FCP port-to-port connectivity is software or hardware enforced (depending on the director or switch firmware release level).
  - FICON architecture controls connectivity through a host-based HCD program, a director or switch-resident management server called the control unit port (CUP), and a director or switch-resident prohibit dynamic connectivity mask (PDCM) array associated with each logical port address. The CUP and PDCM arrays support hardware enforcement of connectivity control to all port connections; therefore, when a director or switch is set to FICON management style, zoning information is restricted by the hardware instead of by the name server.
  - When the director or switch is set to Open Systems management style, CUP support and the PDCM array are disabled. For FICON devices attached to the director or switch, the user must manage connectivity to match logical port addressing established through the host-based HCD program. For example, if a FICON host expects connectivity through logical port address **1C**, the user must ensure the host is connected to physical port number **24**. Refer to Figure 46 for the physical port number and logical port address map.
- The FCP protocol supports multiple domains (multiswitch fabrics), while the FICON protocol may or may not be limited to a single domain (single-switch fabrics), depending on the director or switch firmware release level as follows:
  - For earlier versions of director or switch firmware (prior to version 4.0), the FICON protocol is limited to a single domain (single-switch fabric) due to single-byte Fibre Channel link address limitations inherited from ESCON. Consequently, when a director or switch is set to FICON management style (FICON compliant), E_Port connections (ISLs) are not

allowed with another fabric switch. The director or switch reports an attempted E_Port connection as invalid and prevents the port from coming online.

— For later versions of director or switch firmware (version 4.0 and later), the domain field of the destination ID is added to the Fibre Channel link address, thus specifying the link address on source and target fabric elements and enabling E_Port (ISL) connectivity. This connectivity is called FICON cascading. For additional information, refer to "FICON Cascading" on page 124.

- When employing inband (Fibre Channel) director or switch management, the open-systems management server (OSMS) is associated with the FCP protocol, and the FICON management server (FMS) is associated with the FICON protocol. Management server differences tend to complicate security and control issues.

  Each server provides facilities to change zoning information (FCP protocol) or the logical port address-based connectivity configuration (FICON protocol), but neither provides sufficient functionality for both protocols.

Features that Impact Protocol Intermixing

The following features impact how a director or switch behaves when deployed in an intermixed environment:

- Hardware-Enforced Zoning
- SANtegrity Binding
- FICON Cascading

Hardware-Enforced Zoning

Hardware-enforced zoning (hard zoning) allows a user to program director or switch route tables that enable hardware logic to route Fibre Channel frames. This process prevents traffic between source and destination devices not in the same zone. Hard zoning provides the open-systems environment with the same protection that PDCM arrays provide in the FICON environment.

In environments that include discovery-oriented devices (FCP) and definition-oriented devices (FICON), system administrators must keep device definitions and zoning definitions synchronized. Hard zoning enforces zoning information at the director or switch level and ensures that the information takes precedence over access definitions configured at the device level. This provides a security element that is useful for mixed environments that use both definition and discovery. For additional information, refer to "Zoning" on page 154.

SANtegrity Binding

The SANtegrity Binding feature (including both fabric binding and switch binding) allows the creation of reliable SAN configurations and provides a mechanism for attached devices to query the user-configured security level employed in a SAN. The feature significantly reduces the impacts of accidental or operator-induced errors.

Fabric binding defines the directors and switches allowed to participate in a fabric, thus preventing accidental fabric merges. Switch binding defines the devices allowed to connect to directors and switches in a fabric, thus providing additional security in SAN environments that must manage a large number of devices. For additional information, refer to "SANtegrity Binding" on page 148.

FICON Cascading

FICON is most often deployed in SANs that have high data integrity and reliability standards. However, the initial FICON architecture was limited to one domain (i.e., a single-switch fabric), which creates severe distance and connectivity limitations. These data standards and the requirement for FICON fabrics in SANs led to protocol changes that support FICON cascading.

FICON cascading allows an IBM eServer zSeries processor to communicate with other zSeries processors or peripheral devices (such as disks, tape libraries, or printers) through a fabric consisting of two or more FICON directors or switches. Cascaded FICON fabrics also provide high end-to-end data integrity to ensure changes to a data stream are always detected and rectified, and that data is always delivered to the correct fabric end point. For additional information, refer to "FICON Cascading" on page 124.

A related feature to consider is the announcement of FCP support for IBM eServer zSeries processors. This development accelerates the requirement for intermix protocol fabrics, because primary processors now support both FICON and FCP.

Protocol Intermixing Best Practices

The release of firmware version 6.0 and the related *HAFM* application simplifies the deployment of protocol intermixed SANs. The single-switch operating mode is eliminated, and the Element Manager graphical user interface (GUI) provides an open systems or FICON management style. Users can toggle between management styles with the director or switch online.

However, the firmware and *HAFM* application do not prevent FCP and FICON device configurations that may interfere with each other. A successful intermix environment requires a set of best practice conventions as follows:

1. **Upgrade fabric element firmware to a common version** — Ensure fabric elements are operating at a common firmware level. This reduces errors due to director or switch incompatibility. Firmware Version 4.0 or higher is required to support FICON cascading. Firmware Version 6.0 or higher is recommended.
2. **Upgrade fabric element software to a common version** — Ensure fabric elements are operating at a common software level. This simplifies fabric fault isolation and reduces errors due to director or switch incompatibility.
   - When a director or switch is set to open systems management style, a traditional Fibre Channel fabric is supported. Inband management through the FMS or OSMS is also supported. The key concern is to avoid disrupting installed FCP devices when connecting FICON devices to a fabric element, and modifying configurations to facilitate FICON communication. The *Element Manager* application does not use logical port addressing or display the **Configure Allow/Prohibit Matrix - Active** dialog box. A PDCM array is not supported, and the HCD defined by an attached host describes FICON connectivity requirements.
   - When a director or switch is set to FICON management style, either multiple domains (fabric elements) are supported, or only a single domain (fabric element) is supported, depending on the firmware release level. Inband management through the FMS or OSMS is also supported. The *Element Manager* application for a PDCM array is configured at the **Configure Allow/ Prohibit Matrix - Active** dialog box. The array activates all or a subset of the connectivity paths established by a host-based HCD.

     When using firmware prior to version 4.0 and the FICON management style, ports are set to F_Port operation, thus eliminating E_Port capability (ISL and fabric capability).
   - When using inband director or switch management, either (or both) of the FMS or OSMS features can be enabled. When either (or both) features are enabled, the director or switch can be set to open systems or FICON management style.
3. **Upgrade fabric elements to a common feature set** — Ensure that a common set of PFE-keyed optional features (refer to "Optional Features" on page 163) are installed on each fabric element. This reduces errors due to

director or switch incompatibility. In addition, the SANtegrity Binding feature (with **Enterprise Fabric Mode** enabled) is required to support FICON cascading.

4. **Logically assign ports** — To organize devices into manageable groups for zoning, director or switch ports should be logically assigned to FCP port groups and FICON port groups. Although FICON devices can be zoned by device WWN, they must also be assigned logical port addresses that correspond to the port addresses configured by the attached host HCD. FICON devices must be attached to these assigned ports. In addition, PDCM arrays affect port connections at the hardware level, so a range of port addresses must be established for FCP device use, and a separate range of port addresses must be established for FICON device use. FCP ports should always be configured to allow communication with each other but disallow communication with FICON ports, and vice versa.
5. **Configure FICON cascading** — Configure and enable FICON cascading for all fabric elements. Refer to "FICON Cascading Best Practices" on page 126 for instructions. As part of this step, ensure that the SANtegrity Binding feature key is installed and **Enterprise Fabric Mode** is enabled for all directors and switches.
   — In conjunction with the SANtegrity Binding feature (fabric and switch binding), consider enabling port binding from a director or switch's *Element Manager* application. Port binding explicitly defines (by WWN or nickname) the device allowed to attach to a Fibre Channel port and provides additional security when logically allocating ports to FCP and FICON groups. Although this process creates additional configuration overhead, port binding is useful for implementations that require protection from accidental misconfigurations.
6. **Configure PDCM arrays** — For each director or switch managed by the FICON management style, define the allow and prohibit settings for FICON device connectivity. Use the *Element Manager* application's **Configure Allow/Prohibit Matrix - Active** dialog box. Port connectivity assignment (step 4) should be reflected in PDCM arrays for FICON connectivity management. The baseline configuration for each fabric element must prohibit communication between FICON and FCP devices.
   — Because PDCM arrays affect port connections at the hardware level, it is imperative to establish a range of port addresses for FCP use and another range for FICON use. FCP-assigned ports should be configured to allow communication with each other and prohibit communication with FICON-assigned ports, and vice versa.

- — On the **Configure Allow/Prohibit Matrix - Active** dialog box, assigning port names to logical port addresses is another practice that should be followed. For example, the port name for all FCP devices could begin with FCP or OS to indicate the associated port addresses attached to open-systems devices. This information emphasize which ports are FCP ports and which are FICON ports, and gives a user the ability to better manage the connectivity matrix.
- — Caution should be exercised when using a PDCM array to prohibit E_Port connectivity. For additional information, refer to "PDCM Arrays" on page 149.

7. **Configure zoning** — Well-behaved intermix environments require the creation of separate zones for FCP and FICON devices. Group all FICON devices into one zone, and then group FCP devices into multiple zones in traditional fashion to facilitate typical open-systems communication.
   - — Be aware that FICON devices do not use the Fibre Channel name server, therefore name server-based zoning does not affect FICON connectivity. However, the name server does affect distribution of registered state change notification (RSCN) service requests to FICON devices. If a FICON device is not in the same zone as other devices, state changes are not properly communicated.
   - — All FICON devices must be included in the same zone to facilitate proper state change notification. This is achieved by creating a unique FICON zone or using the default zone. It is best to disable the default zone and explicitly create a unique zone for all FICON devices. Regardless of the director or switch operating mode, FCP devices must be zoned in the traditional fashion, and FICON devices must be zoned to provide isolation from the FCP devices. All FICON devices must be included in the same zone to facilitate proper state change communication.
   - — When establishing a zoning configuration, FICON devices must be assigned to director or switch port addresses that correspond to port HCD-assigned address definitions configured by the attached host. Associated FICON devices must be connected to the ports as configured.
   - — Note the reciprocal nature of zoning configurations and PDCM arrays. When configuring zoning, all FICON devices are placed in one zone and FCP devices are zoned normally. When configuring definitions in a PDCM array, all FCP devices are configured to allow communication only with each other, and FICON devices are configured normally.

— FICON port addressing provides the ability to swap ports for maintenance. In general, swapping ports in intermix environments does not affect the practices described. However, if a user implements zoning using a domain ID and port numbers, zoning information must be updated contiguous with the port swap operation.

Multiple Data Transmission Speeds in a Single Fabric

The Director 2/64, Edge Switch 2/16, and Edge Switch 2/32 support auto-sensing of 1.0625 and 2.125 Gbps device connections. The introduction of a higher data transmission speed to the SAN design provides several benefits and alternatives:

- **High-speed device connectivity** — As Fibre Channel devices and HBAs evolve and become 2.125 Gbps-capable, higher-speed switches are required to provide basic fabric connectivity.
- **Better fabric performance** — As a connection between edge switches, a 2.125 Gbps ISL delivers double the bandwidth of a 1.0625 ISL. Fibre Channel devices that are not 2.125 Gbps-capable benefit from a higher-speed ISL, because 1.0625 Gbps traffic is multiplexed and transmitted through the 2.125 Gbps ISL.
- **Additional port count** — If additional ISL bandwidth is not required for fabric performance, 2.125 Gbps connectivity allows the number of ISL connections to be reduced, thus yielding additional director or switch ports for device connectivity.

When installing 2.125 Gbps-capable fabric elements in a core-to-edge topology, deploy the directors or switches at the fabric core to provide end-to-end high-speed ISL capability. If 2.125 Gbps device connectivity is required, attach the devices to the core director or switch as Tier 1 devices. If possible, employ device locality by connecting 2.125 Gbps devices to the same director or switch.

Fibre Channel Distance Extension

Connectivity requirements for a SAN differ from the requirements for a data network such as a LAN, MAN, or WAN. These differences are summarized as follows:

- Data networks (LANs, MANs, and WANs) usually offer best-effort communication service, relying on upper-level protocols for end-to-end transport. SANs require high-reliability communications and are intolerant of data loss or retransmission.

- Data networks introduce variable delay and usually support high latency. SANs require minimal delay and latency.
- Data networks rely on a software protocol stack such as Transmission Control Protocol/Internet Protocol (TCP/IP) to provide communications. Such stacks impose prohibitive performance penalties in SANs because data traffic quickly overloads servers.

Because of these differences, SANs are based on Fibre Channel technology optimized for storage environments and offer high-speed, low-overhead communication between servers and storage devices. Data networks are often implemented using Internet Protocol (IP) over gigabit Ethernet. IP is appropriate for data networking because a high level of protocol processing is provided. The protocol conversion approaches to integrating Fibre Channel fabric SANs over a geographically dispersed network (WAN extension) are:

- Fibre Channel over TCP/IP (FCIP).
- Internet Fibre Channel Protocol (iFCP)
- Internet Small Computer Systems Interface Protocol (iSCSI).

FCIP Protocol

The FCIP protocol encapsulates Fibre Channel frames (Fibre Channel or SCSI protocol) into IP packets and fabric domains to IP addresses. This process of encapsulating one information packet inside another is called protocol tunneling. With FCIP, a single SAN fabric is created by connecting multiple SAN islands through IP network tunnels. Figure 47 illustrates FCIP WAN extension.

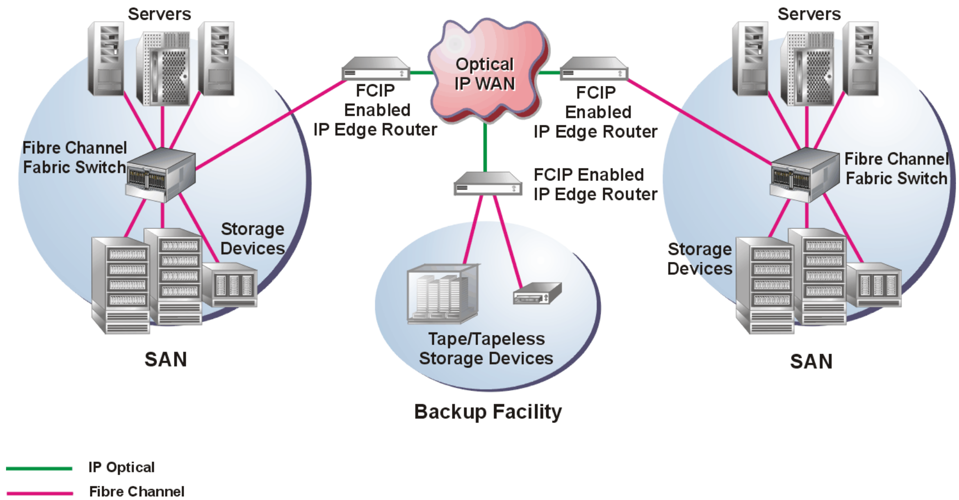

Figure 47: FCIP WAN Extension

FCIP supports existing Fibre Channel SAN hardware and software, while allowing SAN-connected data to be accessed over an IP network infrastructure. FCIP allows data to be accessed remotely without altering the SAN fabric and maintains critically valuable bandwidth, data integrity, and flow control.

Applications appropriate for FCIP include storage-to-storage operations such as extended Fibre Channel SAN interconnection, data protection, outsourced storage services, content distribution, and centralized management of distributed resources.

Planning and implementation FCIP requires available director or switch ports at the fabric core and installation of an edge switch (Fibre Channel-to-IP gateway) between the SAN fabric and IP network. A SAN director or switch port communicates with a remote director or switch port through a protocol tunnel established by the Fibre Channel-to-IP gateway installed at each end of the TCP/IP network. The fabric SAN is extended across the IP network, yet Fibre Channel servers, storage devices, and software are not altered.

iFCP Protocol

iFCP is a TCP/IP-based protocol for connecting distributed Fibre Channel SANs using an IP infrastructure in place of Fibre Channel switching and routing elements. The protocol differs from FCIP in that iFCP is a gateway-to-gateway architecture, while FCIP specifies a protocol tunnel between SAN islands. With

iFCP, each connected SAN fabric is maintained separately from the others, while the IP network provides connectivity, congestion control, error detection, and error recovery. Figure 48 illustrates iFCP WAN extension.

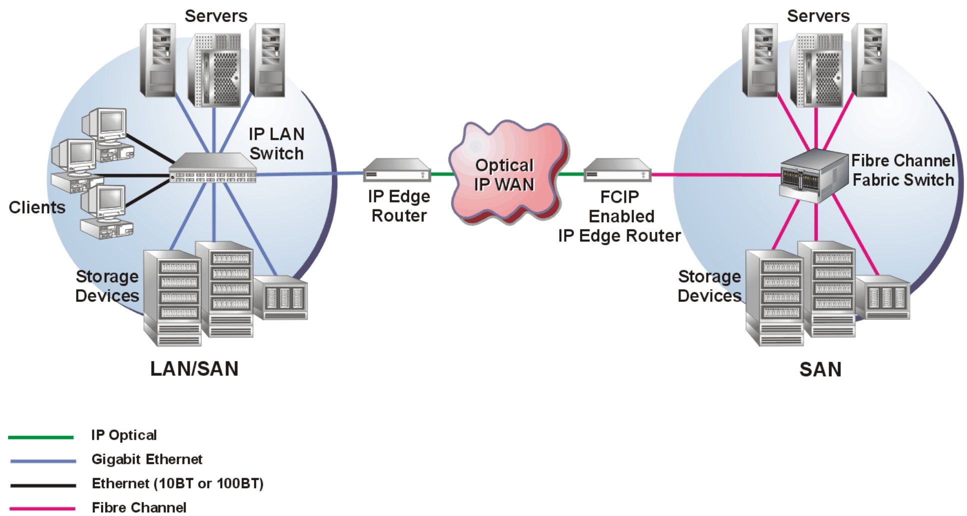

Figure 48: iFCP WAN Extension

iSCSI Protocol

iSCSI is a TCP/IP-based protocol for establishing and managing connections between IP-based storage devices, hosts, and clients. iSCSI operates on top of TCP, moving block data (iSCSI packets) over an IP Ethernet network. This protocol consolidates SANs into a single IP network for data and storage traffic, using Ethernet (not Fibre Channel) for the SAN and IP for the WAN. iSCSI requires equipping both server and storage systems with iSCSI components, resulting in additional capital costs and higher server overhead associated with TCP processing. Applications include storage access where performance is not critical. Figure 49 illustrates iSCSI WAN extension.

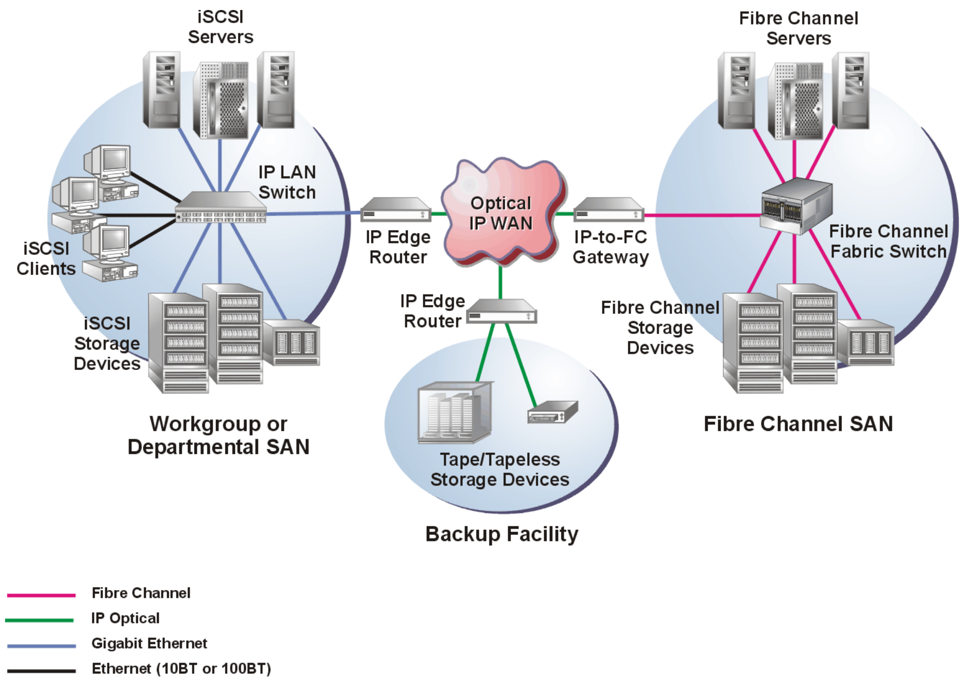

Figure 49: iSCSI WAN Extension

FICON Cascading

The initial FICON architecture did not permit connection of multiple directors or switches, because the protocol specified a single byte for the link (port) address definition in the input-output configuration program (IOCP). The link address only defined the Port_ID for a unique domain (director or switch).

The current FICON architecture provides two-byte addressing that allows the IOCP to specify link (port) addresses for any number of domains by including the domain address with the Port_ID. FICON fabrics can now be configured using multiple director and switches (FICON cascading). In a cascaded FICON environment, at least three Fibre Channel links are involved:

- The first link is between the FICON channel card (N_Port) of an IBM eServer zSeries processor and a director or switch F_Port.
- The second link is an ISL between two director or switch E_Ports.
- The final link is from a director or switch F_Port to a FICON adapter card (control unit N_Port) in a storage device, tape device, or other peripheral.

These Fibre Channel links connect FICON fabric elements and provide a physical transmission path between a channel and control unit. Users may configure multiple ISLs between cascaded FICON directors or switches to ensure redundancy and adequate bandwidth.

High-Integrity Fabrics

Cascaded FICON directors and switches must support high-integrity fabrics. Fabric elements must have the SANtegrity Binding feature installed and operational with **Enterprise Fabric Mode** enabled. High-integrity fabric architecture support includes:

- **Fabric binding** — Only directors or switches with fabric binding installed are allowed to attach to specified fabrics in a SAN. Specifically:
  - Fabric elements without a SANtegrity Binding feature key are prohibited from connecting to fabric elements with an active SANtegrity Binding feature key.
  - Inherent to directors and switches with an active SANtegrity Binding feature key is a fabric membership list (comprised of acceptable WWNs and domain IDs) of the elements logged into the fabric. This membership list is exchanged between fabric elements, and an element with an

incompatible list is isolated from the fabric. Membership list data eliminates duplicate domain IDs and other address conflicts and ensures a consistent, unified behavior across the fabric.

- **Switch binding** — Switch binding allows only specified devices and fabric elements to connect to specified director or switch ports.
- **Insistent domain ID** — When enabled through the **Enterprise Fabric Mode** dialog box, this parameter ensures duplicate domain IDs are not used within a fabric. It also ensures a fabric element cannot automatically change its domain ID when a director or switch with a duplicate domain ID attempts to join the fabric. The invalid (duplicate domain ID) fabric element is rejected, and intentional user intervention is required to change the domain ID to a valid number.

For additional information about the SANtegrity Binding feature, refer to "SANtegrity Binding" on page 148.

Minimum Requirements

The following are minimum hardware, firmware, and software requirements to configure and enable a FICON-cascaded SAN:

- A single-vendor switching environment with two or more of the following Directors or Edge Switches:
  - Director 2/64 or 2/140 .
  - Edge Switch 2/16 or 2/32.
- Enterprise Operating System (E/OS) firmware version 4.0 or later must be installed on all directors or switches. E/OS firmware version 6.0 is recommended. All fabric elements must be at the same firmware version level.
- The SANtegrity Binding feature key must be installed and enabled on all directors and switches. **Enterprise Fabric Mode** must also be enabled on all fabric elements.
- HAFM Version 6.3 or later must be installed on the director or switch management server. HAFM Version 8.02 is recommended.
- One or more of the following IBM servers with FICON or FICON Express™ channel adapter cards:
  - eServer zSeries 800 (z800) processor.
  - eServer zSeries 900 (z900) processor.

— eServer zSeries 990 (z990) processor.

Note: FICON cascading is not supported for IBM S/390 Parallel Enterprise Servers (Generation 5 or Generation 6).

- The z/OS version 1.3 or version 1.4 operating system (with service as defined in PSP Buckets for device type 2032, 2042, 2064, or 2066) must be installed on the IBM server.
- Licensed Internal Code (LIC) driver 3G at microcode level (MCL) J11206 or later must be installed on the IBM server.

FICON Cascading Best Practices

A successful FICON-cascaded SAN environment requires a set of best practice conventions as follows:

1. **Connect fabric elements** — Establish one or more ISLs between cascaded directors of fabric switches as follows:
   a. Ensure fabric elements are defined in the SAN management application. If the elements must be defined, refer to the appropriate switch or director installation manual for instructions.
   b. Ensure the preferred domain ID for each director or switch is unique and does not conflict with the ID of another fabric element.
   c. Ensure the R_A_TOV and E_D_TOV values for fabric elements are identical.
   d. Route multimode or singlemode fiber-optic cables (depending on the type of transceiver installed) between customer-specified E_Ports at each fabric element.
2. **Verify operation of local FICON applications** — Ensure the ISL connection(s) do not disrupt fabric element operation nor disrupt local FICON (non-cascaded) traffic. Perform this step at each director or switch.
   a. At the SAN management application's physical map, right-click the director or switch product icon, then select **Element Manager** from the pop-up menu. The *Element Manager* application opens.

b. If required, click the **Hardware** tab. The **Hardware View** (Figure 19) displays. Verify that the status bar at the bottom left corner of the window displays a green circle, indicating director or switch status is operational. If a problem is indicated, go to **MAP 0000: Start MAP** in the product-specific *Installation and Service Manual*.

c. Verify operation of non-cascaded FICON applications at each director or switch.

3. **Verify ISL operation** — Ensure ISL connectivity between fabric elements. Perform this step at each director or switch.

a. In the *Element Manager* application's **Hardware View**, double-click the graphical E_Port connector used for the ISL. The **Port Properties** dialog box displays (Figure 50).

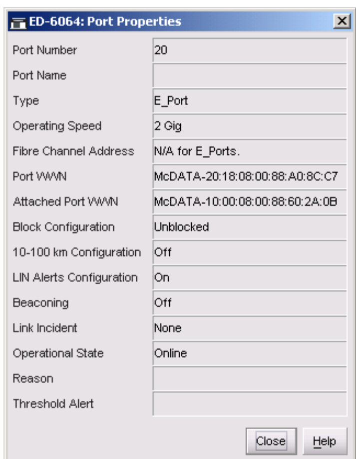

Figure 50: Port Properties Dialog Box

b. Ensure that the **Link Incident** field displays **None** and the **Reason** field is blank. If an ISL segmentation or other problem is indicated, refer to the diagnostics information described in the appropriate service manual for your Director or Edge Switch.

c. Click **Close** to close the dialog box and return to the **Hardware View**.

4. **Install SANtegrity Binding on fabric elements** — Configure and enable the SANtegrity Binding feature at each director or switch as follows:

a. In the *Element Manager* application, install the SANtegrity Binding PFE key. Refer to the installation instructions in the appropriate installation manual for your Director or Edge Switch.

b. In the *HAFM* application, configure fabric binding. Refer to installation instructions in the *HA-Fabric Manager User Guide.*

c. In the *Element Manager* application, configure switch binding. Refer to the installation instructions in the appropriate installation manual for your Director or Edge Switch.

5. **Ensure FICON devices are logged in** — Verify FICON devices are logged in to each director or switch as follows:

   a. At the *Element Manager* application's **Hardware View**, click the **Node List** tab. The **Node List View** displays (Figure 51).

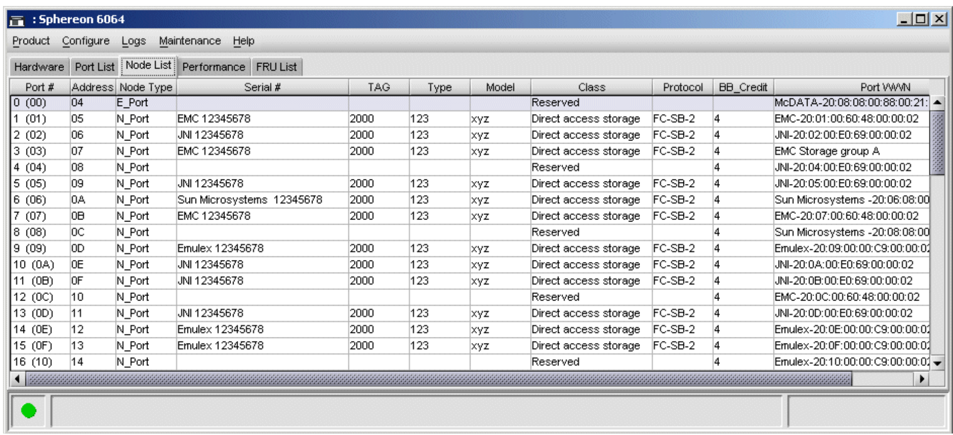

Port #	Address	Node Type	Serial #	TAG	Type	Model	Class	Protocol	BB_Credit	Port WWN
0 (00)	04	E_Port					Reserved			McDATA-20:08:08:00:88:00:21:
1 (01)	05	N_Port	EMC 12345678	2000	123	xyz	Direct access storage	FC-SB-2	4	EMC-20:01:00:60:48:00:00:02
2 (02)	06	N_Port	JNI 12345678	2000	123	xyz	Direct access storage	FC-SB-2	4	JNI-20:02:00:E0:69:00:00:02
3 (03)	07	N_Port	EMC 12345678	2000	123	xyz	Direct access storage	FC-SB-2	4	EMC Storage group A
4 (04)	08	N_Port					Reserved		4	JNI-20:04:00:E0:69:00:00:02
5 (05)	09	N_Port	JNI 12345678	2000	123	xyz	Direct access storage	FC-SB-2	4	JNI-20:05:00:E0:69:00:00:02
6 (06)	0A	N_Port	Sun Microsystems 12345678	2000	123	xyz	Direct access storage	FC-SB-2	4	Sun Microsystems -20:06:08:00
7 (07)	0B	N_Port	EMC 12345678	2000	123	xyz	Direct access storage	FC-SB-2	4	EMC-20:07:00:60:48:00:00:02
8 (08)	0C	N_Port					Reserved		4	Sun Microsystems -20:08:08:00
9 (09)	0D	N_Port	Emulex 12345678	2000	123	xyz	Direct access storage	FC-SB-2	4	Emulex-20:09:00:00:C9:00:00:0
10 (0A)	0E	N_Port	JNI 12345678	2000	123	xyz	Direct access storage	FC-SB-2	4	JNI-20:0A:00:E0:69:00:00:02
11 (0B)	0F	N_Port	JNI 12345678	2000	123	xyz	Direct access storage	FC-SB-2	4	JNI-20:0B:00:E0:69:00:00:02
12 (0C)	10	N_Port					Reserved		4	EMC-20:0C:00:60:48:00:00:02
13 (0D)	11	N_Port	JNI 12345678	2000	123	xyz	Direct access storage	FC-SB-2	4	JNI-20:0D:00:E0:69:00:00:02
14 (0E)	12	N_Port	Emulex 12345678	2000	123	xyz	Direct access storage	FC-SB-2	4	Emulex-20:0E:00:00:C9:00:00:0
15 (0F)	13	N_Port	Emulex 12345678	2000	123	xyz	Direct access storage	FC-SB-2	4	Emulex-20:0F:00:00:C9:00:00:0
16 (10)	14	N_Port					Reserved		4	Emulex-20:10:00:00:C9:00:00:0

Figure 51: Node List View

   b. Inspect the node descriptors and verify that the correct FICON devices (channels and control units) are logged in to each director or switch.

6. **Enable Enterprise Fabric Mode** — Enable **Enterprise Fabric Mode** as follows:

   a. Minimize the *Element Manager* application to display the SAN management application, then select **Enterprise Fabric Mode** from the **Configure** menu. The **Enterprise Fabric Mode** dialog box displays (Figure 52).

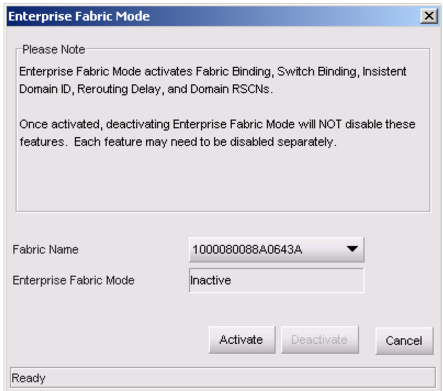

Figure 52: Enterprise Fabric Mode Dialog Box

b. Select the fabric to be configured from the *Fabric Name* drop-down list. The selected fabric's status displays in the *Enterprise Fabric Mode* field.

c. Click **Activate** to close the dialog box and enable **Enterprise Fabric Mode** for the selected fabric.

7. **Verify FICON devices are still logged in** — Maximize the *Element Manager* application. Inspect the **Node List View** (Figure 51) and verify FICON devices (channels and control units inspected in step 5) are still logged in to each director or switch.

8. **Change switch binding enforcement if required** — If the SAN environment is volatile (characterized by a high volume of optical cable connects, disconnects, and movement), change switch binding enforcement to restrict E_Ports only.

   a. In the *Element Manager* application, click the **Hardware** tab. In the **Hardware View**, select **Switch Binding**, and then select **Change State** from the **Configure** menu. The **Switch Binding - State Change** dialog box displays (Figure 53).

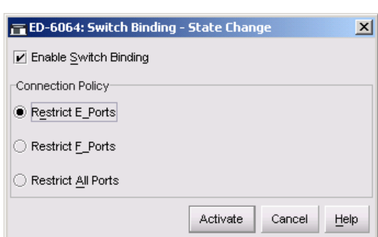

Figure 53: Switch Binding - State Change Dialog Box

b. Ensure that the **Enable Switch Binding** check box is checked (enabled).

c. Select the **Restrict E_Ports** radio button to restrict connections from specific fabric elements to E_Ports. WWNs can be added to the membership list to allow fabric element connection and removed from the list to prohibit fabric element connection. Devices are allowed to connect to any F_Port or FL_Port without restriction.

d. Click **Activate** to close the dialog box and enforce the connection policy.

9. **Update channel path and control unit definitions** — A cascaded FICON environment requires channel entry switch and link address updates to the input/output configuration program (IOCP) as follows:

   a. In the IOCP, define an entry switch ID in the **SWITCH** keyword of the channel path identifier (CHPID) definition.

Note: An entry switch is a fabric director or switch connected to the FICON channel of a zSeries processor and a second fabric director or switch.

---

   b. In the IOCP, define a 2-byte link address (consisting of a switch (or domain) address and port address) for the cascaded switch in the **LINK** keyword of the control unit (**CNTLUNIT**) definition.

Note: An cascaded switch is a fabric director or switch connected to a destination control unit and an entry switch.

---

   c. Run the IOCP to create an input/output configuration data set (IOCDS). The switch ID (CHPID macroinstructions) and 2-byte link address (control unit macroinstructions) are updated in the IOCDS.

   Refer to the IBM *FICON Native Implementation and Reference Guide* (SG24-6266) for additional information.

10. **Verify FICON devices log back in** — Inspect the **Node List View** (Figure 51) and verify FICON devices (channels and control units inspected in step 5) log back in to the fabric as expected.

11. **Verify cascaded FICON operation** — Verify operation of established logical FICON paths between channels and control units, and verify that cascaded FICON traffic is transmitted through the fabric as expected.

Physical Planning Considerations

4

This chapter describes the physical planning considerations for incorporating Hewlett-Packard (HP) Director 2/64s, Director 2/140s, Edge Switch 2/12s, Edge Switch 2/16s, Edge Switch 2/24s, and Edge Switch 2/32s into storage area networks (SANs) and Fibre Channel fabric topologies. This chapter provides planning considerations and recommendations for:

Port Connectivity and Fiber-Optic Cabling

This section provides planning recommendations for director and switch port connectivity and fiber-optic cabling. Recommendations are provided for:

- Port Requirements
- Optical Transceivers
- Extended-Distance Ports
- High-Availability Considerations
- Cables and Connectors
- Routing Fiber-Optic Cables

Port Requirements

Plan for sufficient shortwave laser ports and longwave laser ports to meet the needs of the SAN configuration. The number of ports required is equal to the number of device connections (including redundant connections), plus the number of interswitch links (ISLs) between fabric elements, plus the total number of spare port connections.

WARNING: Director and switch non-open fiber control (non-OFC) laser transceivers are designed and certified for use only with fiber-optic cables and connectors with characteristics specified by HP. Use of other connectors or optical fiber can result in emission of laser power levels capable of producing injury to the eye if viewed directly. Use of non-specified connectors or optical fiber can violate the Class 1 laser classification.

The number of Fibre Channel ports and port operation for directors and switches are described as follows:

- **Director 2/64** — The director is configured from a minimum of 8 universal port module (UPM) cards (32 ports total) to a maximum of 16 UPM cards (64 ports total).

  UPM cards provide four 2.125 Gbps port connections and can be configured with shortwave transceivers, longwave transceivers, and extended longwave transceivers or a combination of all three.

- **Director 2/140** — The director is configured from a minimum of 16 universal port module (UPM) cards (64 ports total) to a maximum of 35 UPM cards (140 ports total).

  UPM cards provide four 2.125 Gbps port connections and can be configured with shortwave transceivers, longwave transceivers, and extended longwave transceivers or a combination of all three.
- **Edge Switch 2/12** — The switch provides up to 12 duplex SFP fiber-optic port transceivers (2.125 Gbps operation only). Shortwave laser, longwave laser, and extended longwave transceivers are available.
- **Edge Switch 2/16** — The switch provides up to 16 duplex small form factor pluggable (SFP) fiber-optic port transceivers (2.125 Gbps operation only). Shortwave laser, longwave laser, and extended longwave transceivers are available.
- **Edge Switch 2/24** — The switch provides up to 24 duplex SFP fiber-optic port transceivers (2.125 Gbps operation only). Shortwave laser, longwave laser, and extended longwave transceivers are available.
- **Edge Switch 2/32** — The switch provides up to 32 duplex SFP fiber-optic port transceivers (2.125 Gbps operation only). Shortwave laser, longwave laser, and extended longwave transceivers are available.

Optical Transceivers

Shortwave optical transceivers provide a connection for multimode cable with a core diameter of 50 microns and a cladding diameter of 125 microns (50/125) or multimode cable with a core diameter of 62.50 microns and a cladding diameter of 125 microns (62.5/125). A 50/125 micron cable allows a maximum switch-to-device or switch-to-switch distance of up to 300 meters at a 2.125 Gbps data transmission speed. A 62.5/125 micron cable allows a maximum switch-to-device or switch-to-switch distance of up to 150 meters at a 2.125 Gbps data transmission speed.

A 62.5 micron cable is supported only for the use of existing cable plants. HP recommends the use of 50 micron cables for new installations.

Longwave optical transceivers provide a connection for single-mode cable with a core diameter of 9 microns and a cladding diameter of 125 microns (9/125). Depending on transceiver type, a 9/125 micron cable allows switch-to-device or switch-to-switch distances of 10, 20, or 35 kilometers.

HP supplies cables for 10 and 35 kilometers. For longer distances, you need Fibre Channel repeaters or wave division multiplexing (WDM) devices.

Consider the following when determining the number and type of each transceiver to use:

- Distance between a director or switch and the attached Fibre Channel device or between fabric elements communicating through an ISL.
- Cost effectiveness.
- Device restrictions or requirements with respect to existing fiber-optic cable (multimode or single-mode).

Data Transmission Distance

Data transmission distance is the primary factor governing the choice of transceiver type and optical fiber. If the transmission distance is:

- Less than 150 meters, multimode or single-mode optical fiber and any type of optical transceiver can be used.
- Between 150 and 300 meters, 50/125-micron multimode or single-mode optical fiber and any type of transceiver can be used.
- Over 300 meters, only single-mode optical fiber and a longwave laser transceiver can be used. A 62.5 micron cable is only supported for the use of existing cable plants. HP recommends the use of 50 micron cables for new installations.

Variables such as the number of patch panel connections, link speed, grade of fiber-optic cable, device restrictions, application restrictions, buffer-to-buffer credit limits, and performance requirements can affect transmission distance.

Cost-Effectiveness

Cost is another factor governing the choice of transceiver type and optical fiber. Shortwave laser transceivers and multimode cable offer a less expensive solution if data transmission distance is not critical.

Device or Cable Restrictions

The choice of transceiver and cable type may be restricted or dictated by:

- **Device restrictions** — Some devices may be restricted to use of only one type of transceiver (shortwave or longwave). Refer to the supporting documentation delivered with the product for information.
- **Existing cable restrictions** — The enterprise may contain only one type of fiber-optic cable (multimode or single-mode) and the customer may be required to use the existing cables.

Extended-Distance Ports

Through longwave laser transceivers and repeaters or dense wavelength division multiplexing (DWDM) equipment, Directors and Edge Switches support Fibre Channel data transmission distances of up to 100 km at 1 Gbps, or 50 km at 2 Gbps. The extended distance feature is enabled on a port-by-port basis by activating the **10-100 km** check box for a specified port on the *Element Manager* application's **Configure Ports** dialog box. This feature provides extended distance support using Fibre Channel protocol only and does not support distance extension using Fibre Channel over Internet Protocol (FCIP) conversion.

When a port is configured for extended distance operation, the buffer-to-buffer credit (BB_Credit) value for the port is automatically set to 60. This value provides sufficient buffering to handle frame processing for link distances up to 100 km. When a director or switch port is configured to support extended link distances, the attached device (or attached fabric element) must also support extended distance operation and be configured to use a higher BB_Credit value to maintain link efficiency.

If the extended distance feature is enabled for a port that is not installed or does not support extended distance operation, the configuration for the feature is ignored. In addition, a director or switch port configured for extended distance operation cannot transmit broadcast frames to other ports in a Fibre Channel fabric.

High-Availability Considerations

To provide high device availability, critical servers, storage devices, or applications should be connected to more than one fabric element (director or switch) or to more than one fabric. To determine if dual-connection capability exists for a device, refer to the associated device documentation. To provide high fabric availability, consider the use of multiple fabric elements, multiple ISLs, or redundant fabrics. Refer to "Fabric Availability" on page 105 for additional information.

Plan to maintain unused (spare) director and switch ports if port connections must be quickly moved and re-established after a failure. If an individual port or an entire port card fails, optical transceivers or port cards can be removed and replaced, spare port connections identified (through the *Element Manager* application), and fiber-optic cables rerouted and reconnected while the director or switch is operational.

Cables and Connectors

This section provides Fibre Channel cable and connector planning information as follows:

- Cables for all directors and switches.
- SFP transceivers for Director 2/64, Director 2/140, Edge Switch 2/12, Edge Switch 2/16, Switch 2/24, and Edge Switch 2/32.

Cables

Fiber-optic jumper cables are required to connect director and switch ports to servers, devices, distribution panels, or other elements in a multi-switch fabric. Depending on the attached device, director port, or switch port, use one of the following types of cable:

- Graded-index multimode cable with a core diameter of 50 microns and a cladding diameter of 125 micron (50/125). The cable provides a transmission distance of up to 300 meters at 2.125 Gbps and connects to shortwave ports that transmit light at an 850 nanometer (nm) wavelength. The cable typically has an orange jacket.
- Graded-index multimode cable with a core diameter of 62.50 microns and a cladding diameter of 125 microns (62.5/125). The cable provides a transmission distance of up to 150 meters at 2.125 Gbps and connects to shortwave ports that transmit light at an 850 nm wavelength. The cable typically has an orange jacket.
- Dispersion-unshifted (step-index) single-mode cable with a core diameter of nine microns and a cladding diameter of 125 microns (9/125). Depending on transceiver type, the cable provides a transmission distance of up to 10, 20, or 35 kilometers and connects to longwave ports that transmit light at a 1310 nm wavelength. The cable typically has a yellow jacket.

  HP supplies cables for 10 and 35 kilometers. For longer distances, you need Fibre Channel repeaters or wave division multiplexing (WDM) devices.

Director and Switch Connectors

Multimode or single-mode cables attach to Director 2/64, Director 2/140, Edge Switch 2/12, Edge Switch 2/16, Edge Switch 2/24, and Edge Switch 2/32 ports with SFP transceivers with LC duplex connectors. Figure 54 illustrates an SFP transceiver and LC duplex connector.

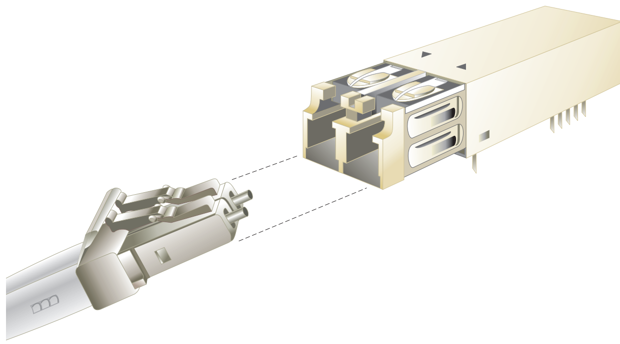

Figure 54: SFP transceiver and LC duplex connector

Routing Fiber-Optic Cables

Follow a logical plan for routing fiber-optic cables to avoid confusing connections during installation and operation. Route cables from the access holes at the bottom or top of the equipment rack, and then to director and switch ports.

Leave enough slack in the cables to allow cable movement for UPM card or optical transceiver removal and replacement or possible rerouting of the cable to another port.

When routing fiber-optic cables and estimating cable lengths, consider:

- Cable routing inside the equipment rack to different port locations and installation position of the director or switch (top or bottom of the rack). Plan for 1.0 meter (39.37 inches) of extra cable for routing through restraint mechanisms and rerouting cables to other ports.
- Cable routing outside the equipment rack. Plan for 1.5 meters (5 feet) of cable outside the rack to provide slack for service clearance, limited rack movement, and inadvertent cable pulls.
- Cabling distance to servers, storage devices, and other fabric elements (for multi-switch fabric support).

The need for additional fiber-optic cabling could grow rapidly. More cables may be required for connections to additional servers or storage devices or for connections to additional fabric elements as a multi-switch fabric is developed. The director or switch may need to be moved for more efficient connection to

other units, while maintaining its original connections. To account for these possibilities, consider installing excess fiber-optic cable, especially in hard-to-reach places like underground trenches.

HAFM Appliance, LAN, and Remote Access Support

Out-of-band (non-Fibre Channel) console access to directors and switches is provided to perform a variety of operations and management functions. These functions are performed from one or more of the following consoles:

- Through the HAFM appliance attached to an Ethernet port on a director control processor (CTP) card or switch front panel.
- Through a remote personal computer (PC) or workstation connected to the HAFM appliance through a customer intranet.
- Through a simple network management protocol (SNMP) management workstation connected through the customer intranet.
- Through a PC with a Web browser and Internet connection to the director or switch through a LAN segment.
- Through a PC with a direct serial connection to the director or switch maintenance port. The maintenance port is used by installation personnel to configure product network addresses.
- Through a PC with a modem connection to the HAFM appliance. The modem is for use by support center personnel only.

HAFM Appliance

The HAFM appliance is mounted in a slide-out drawer in the equipment rack. The server supports up to 48 HP directors or switches (managed products). The server is used to configure products and the associated *HAFM* and *Element Manager* applications, monitor product operation, change configurations, download firmware updates, and initiate diagnostics.

Note: The Edge Switch 2/12 is not supported by the HAFM appliance.

An HAFM appliance failure does not affect port connections or functions of an operational director or switch. The only operating effect of a server failure is loss of remote access, configuration, management, and monitoring functions.

HAFM Appliance Connectivity

The HAFM appliance provides an auto-detecting 10/100 Base-T Ethernet interface that connects to a hub. Each director CTP card or switch front panel also provides an auto-detecting 10/100 Base-T Ethernet interface that connects to a hub. A 12-port hub can be ordered from HP and installed at the top front of the equipment rack.

Although directors provide two Ethernet connections to a hub, only one connection is active at a time. The interface on the backup CTP card remains passive until a failure on the active CTP card occurs, at which point the redundant CTP card becomes active using the same media access control (MAC) address as the original interface.

If an optional private intranet is to be used for LAN connections, an optional Ethernet adapter card (not supplied by HP) can be installed in the Personal Computer Memory Card International Association (PCMCIA) slot in the HAFM appliance to provide a connection to a private LAN segment for dedicated director communication.

The HAFM appliance uses a modem connection for service and support of managed products. The modem provides a dial-in capability that allows HP-authorized service personnel to communicate with the HAFM appliance and operate the *HAFM* and *Element Manager* applications remotely.

The modem is also used to automatically dial out to an authorized support center (to report the occurrence of significant system events) using a call-home feature. The call-home feature is enabled in the *Element Manager* application and configured through the dial-up networking feature of Windows 2000.

For directors and switches installed in some legacy environments, call-home notification requires installation of HP Proactive Service software. This service is offered at no additional charge for subsystems covered under an on-site warranty or on-site storage hardware support contract. To register or order Proactive Service software, contact your HP customer service representative.

Connectivity Planning Considerations

Directors, switches, and the HAFM appliance can be delivered in an HP-supplied equipment rack in accordance with customer specifications. Consider the following Ethernet connectivity issues when:

- **Installing additional rack-mount products** — When installing an additional director or switch, the length of Ethernet cable required to provide LAN connectivity is a function of rack position (top, bottom, or adjacent to the slide-out drawer). Ensure cable lengths provide sufficient cable inside the rack to route to the product's Ethernet ports and to allow service clearance.
- **Interconnecting equipment racks** — To increase the products managed by one HAFM appliance, Ethernet hubs in one or more equipment racks must be connected. Plan for an Ethernet cable length that meets the distance requirement between the racks. In addition, plan for an additional 1.5 meters (5 feet) of cable outside the rack to provide slack for service clearance, limited rack movement, or inadvertent cable pulls. Store extra Ethernet cable in the rack or under the computer room raised floor.
- **Consolidating HAFM appliance operation** — For control and efficiency, all directors and switches in a multi-switch fabric should be managed by one HAFM appliance. When products in two or more racks are joined to form a fabric, the PC environment should be consolidated to one server and one or more clients. Plan for Ethernet cabling to interconnect racks and ensure all directors, switches, and PC platforms participating in the fabric have unique IP addresses.

Remote User Workstations

Customer system administrators determine whether to allow access to directors from remote workstations. If administrators allow remote sessions, they may restrict access to selected workstations by configuring the IP addresses of those workstations through the *HAFM* application. When a remote session is allowed, the remote user has the same rights and permissions as if the session were on the local HAFM appliance. Up to 25 *HAFM* application sessions can be simultaneously active (one local from the HAFM appliance and 24 remote).

Remote workstations must have access to the LAN segment on which the HAFM appliance is installed. Director administrative functions are accessed through the LAN and server.

The LAN interface can be:

- Part of the customer's public 10/100 Mbps LAN segment that provides access to managed directors and switches. This product-to-HAFM appliance Ethernet connection is part of the equipment rack installation and is required. Connection of remote workstations through the hub is optional. This type of network configuration using one Ethernet connection through the HAFM appliance is shown in Figure 55. Director 2/64s are used as an example.

This single Ethernet connection is supported by HP, is Open View-Storage Node Manager (OV-SNM) compatible, and is the recommended configuration for a typical HP installation at a customer site. LAN security is provided by restricting password access and disabling the SNMP agent, Embedded Web Server interface, and command line interface (Telnet access) for each managed director or switch.

Note: The Ethernet adapter in the HAFM appliance provides an auto-detecting 10/100 Mbps connection. Depending on speed restrictions imposed by other LAN-attached devices, the LAN segment that connects the HAFM appliance to managed directors and switches operates at either 10 or 100 Mbps.

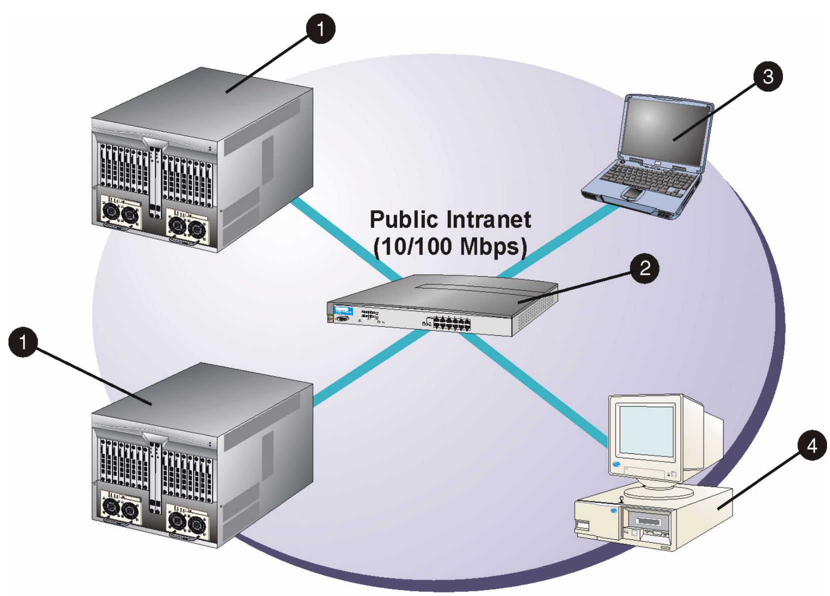

❶ Two Director 2/64s
❷ HP Ethernet hub
❸ HAFM appliance (laptop or rack mount)
❹ Remote user workstation

Figure 55: Typical network configuration (one Ethernet connection)

- Part of a second HAFM appliance interface that connects to the customer's private intranet and allows operation of the *HAFM* and *Element Manager* applications from remote user PCs or workstations. Connection to this LAN

segment is optional and depends on customer requirements. This type of network configuration using both Ethernet connections is shown in Figure 56. Director 2/64s are used as an example.

Although this dual Ethernet connection is supported by HP, it is not OV-SNM compatible, requires installation of an additional PCMCIA LAN adapter card (not supplied by HP), and is not the recommended configuration for a typical new HP installation at a customer site.

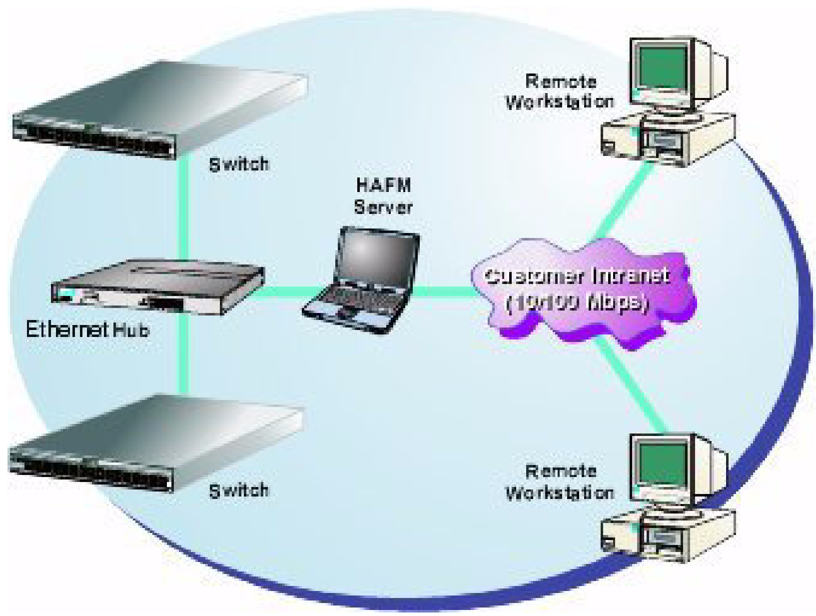

Figure 56: Typical network configuration (two Ethernet connections)

SNMP Management Workstations

An SNMP agent that runs on the HAFM appliance can be configured through the *HAFM* application. This agent implements version 3.1 of the Fibre Alliance management information base (MIB). The agent can be configured to send SNMP trap messages to up to 12 recipients. In addition, there is a separate SNMP agent that runs on each director or switch that is configured through the *Element Manager* application. This agent implements the following MIBs:

- The Fibre Channel Fabric Element MIB (version 3.1).

- A subset of the standard transmission control protocol/internet protocol (TCP/IP) MIB-II definition (RFC1213).
- The director or switch-specific private enterprise MIB.

The director or switch SNMP agent can be configured to send trap messages to up to six recipients. SNMP management is intended only for product monitoring; therefore, the default state of all MIB variables is read-only. When installed on a customer intranet, workstations communicate with directors and switches through the HAFM appliance.

Web Browser Access

The Embedded Web Server interface provides a graphical user interface (GUI) accessed through the Internet (locally or remotely) to manage a single director or switch. If the Embedded Web Server interface is to be implemented:

- Plan for an Internet connection to the LAN segment on which the product is installed. The LAN connection is provided through the customer’s intranet.
- Ensure that adequate security measures are implemented to preclude unauthorized access to managed products. Ensure that IP addresses (uniform resource locators [URLs] for Internet access) of managed products, usernames, and passwords are tightly controlled.

Inband Management Access (Optional)

Inband management console access (through a Fibre Channel port) is provided by enabling user-specified features that allow Open Systems or FICON host control of a director or switch. The features are mutually exclusive; only one can be installed at a time.

Features are enabled through a feature key encoded to work with the serial number of a unique director or switch. A feature key is a case-sensitive alphanumeric string with dashes every four characters.

When the Open Systems management server (OSMS) feature key is enabled at a *Element Manager* application, host control and management of the director or switch is provided through an open-systems interconnection (OSI) device attached to a product port.

When implementing inband product management through an OSI connection, plan for the following minimum host requirements:

- Connectivity to an OSI server with a product-compatible host bus adapter (HBA) that communicates through the Fibre Channel common transport (FC-CT) protocol.
- Installation of a storage network management application on the OSI server. Management applications include Veritas SANPoint Control (version 1.0 or later) or Tivoli NetView (version 6.0 or later).

For information about product-compatible HBAs, third-party SAN management applications, and minimum OSI server specifications, refer to the HP web site.

When the FICON management server (FMS) feature key is enabled in the *Element Manager* application, host control and management of the director or switch is provided through an IBM server attached to a product port. The server communicates with the product through a FICON channel.

When implementing inband product management through a FICON channel, plan for the following minimum host requirements:

- Connectivity to an IBM System/390 (generation 5 or later) or zSeries 900 Parallel Enterprise server with one or more FICON channel cards installed.
- Installation of System Automation for Operating System/390 (SA OS/390) for native FICON, version 1.3 or later, plus service listed in the appropriate preventive service planning (PSP) bucket. The PSP bucket upgrade is HKYSA30.

The minimum OS/390 level for a director or switch without the control unit port (CUP) feature is version 2.6, plus service listed in PSP bucket upgrade 2032, device subset 2032OS390G5+. The minimum OS/390 level for a director or switch with the CUP feature is version 2.1, plus service listed in the preceding PSP bucket for that function.

- A host-attached Hardware Management Console. The console runs the *Hardware Management Console* application (HWMCA) and is the operations and management PC platform for S/390 servers.

For additional information, refer to the IBM publication *System Automation for OS/390, Operations* (GC28-1550).

Security Provisions

Security provisions are available to restrict unauthorized access to a director, switch, or attached Fibre Channel devices. Access to the director or switch (through the *HAFM* application, *Element Manager* application, and Web server interface) is restricted by implementing password protection. Access to attached computing resources (including applications and data) is restricted by implementing one or more of the following security provisions:

- SANtegrity Binding
- PDCM Arrays
- Preferred Path
- Zoning
- Server and Storage-Level Access Control

Password Protection

Access to the *HAFM* and *Element Manager* applications requires configuration of a user name and password. Up to 16 user names and associated passwords can be configured, although only 9 users can log in concurrently (8 remote and 1 local). Each user is assigned rights that allow access to specific sets of product management operations.

Table 3 explains the types of user rights available. A user may have more than one set of user rights granted.

Table 3: Types of User Rights

User Right	Operator Access Allowed
View Only	The user may view product configurations and status but may not make changes. These rights are the default if no other user rights are assigned.
Operator	The operator may view status and configuration information through the *Element Manager* application and perform operational control changes, such as blocking ports and placing the product online or offline.

Table 3: Types of User Rights

User Right	Operator Access Allowed
Product Administrator	The product administrator can make control and configuration changes through the *Element Manager* application.
System Administrator	The system administrator can make control and configuration changes, define users and passwords, and add or remove products through the *HAFM* application.
Maintenance	The maintenance operator can perform product control and configuration changes through the *Element Manager* application and perform diagnostics, maintenance functions, firmware loads, and data collection.

The system administrator can also use the *HAFM* application to assign remote workstation access to directors and switches. Remote sessions can be allowed for anyone on a customer intranet, disallowed completely, or restricted to specific workstations. Remote users must log in to the *HAFM* application with a user name and password, just as when logging in to the local HAFM appliance. Passwords are encrypted when sent across the network. By entering workstation IP addresses at the *HAFM* application, administrators can allow access from all user workstations or only from specific workstations.

For access through the Web server interface, the system administrator provides IP addresses of products to authorized users, assigns access usernames, and controls associated passwords.

SANtegrity Binding

SANtegrity Binding is a feature that enhances data security in large and complex SANs that have numerous fabrics and devices provided by multiple original equipment manufacturers (OEMs), SANs that intermix FCP and FICON protocols, and FICON-cascaded high-integrity SANs. The feature allows or prohibits director or switch attachment to fabrics (fabric binding) and Fibre Channel device attachment to directors or switches (switch binding). The SANtegrity binding feature includes:

- **Fabric binding** — Using the fabric binding feature, an administrator allows only specified directors or switches to attach to specified fabrics in a SAN. This provides security from accidental fabric merges or potential fabric disruption, particularly in environments that use patch panels for centralizing fibers and physical connections. This feature is managed through the *HAFM Manager* application.

- **Switch binding** — Using the switch binding feature, an administrator allows only specified devices and fabric elements to connect to specified director or fabric switch ports. This provides security in environments that include a large number of devices by ensuring that only the intended set of devices attaches to a director or switch. This feature is managed through the *Element Manager* application.

SANtegrity Binding Planning Considerations

Fabric and switch binding enhance data security by controlling and monitoring director, fabric switch, and device connectivity. In fact, installation of the SANtegrity Binding feature is a prerequisite for configuring a high-integrity, FICON-cascaded SAN.

Use of the SANtegrity Binding and zoning features in conjunction with each other must be carefully planned and coordinated. Refer to “Zoning” on page 154 for additional information about zoning.

It is recommended that you obtain planning assistance from the HP professional services organization before implementing the SANtegrity Binding feature with director or switch zoning, especially for multiple fabrics.

PDCM Arrays

PDCM connectivity control is configured and managed at the director or switch level using the **Configure Allow/Prohibit Matrix - Active** dialog box (Figure 57), where the user specifies an array in which logical port addresses are allowed or prohibited from connecting with each other (including E_Port connectivity). To access the dialog box, ensure the FICON management style is enabled for the director or switch, then select the **Allow/Prohibit** and **Active** options from the *Element Manager* application’s **Configure** menu.

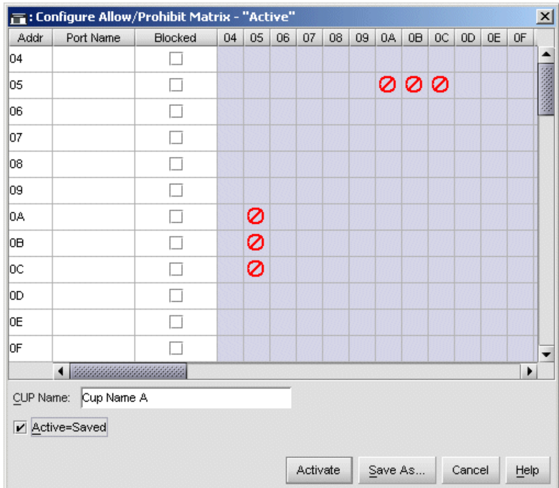

Figure 57: Configure Allow/Prohibit Matrix - Active Dialog Box

Figure 57 shows that port **1** (logical port address **05**) is prohibited from communicating with port **6** (logical port address **0A**), port **7** (logical port address **0B**), and port **8** (logical port address **0C**).

When implementing an array that prohibits E_Port connectivity, be aware that ISLs can be configured as unavailable to attached devices, causing complex routing problems that can be difficult to fault isolate and be incorrectly diagnosed as issues associated with the devices.

As an example of such a problem, refer to the simple two-director fabric illustrated in Figure 58. As shown in the figure, ISL **1** connects Director **A** and Director **B** through logical port addresses **09** and **1A**. ISL **2** connects the directors through logical port addresses **0A** and **1B**. A source server attaches to Director **A** through logical port address **05**. Two destination devices attach to Director **B** through logical port addresses **2C** and **2D**.

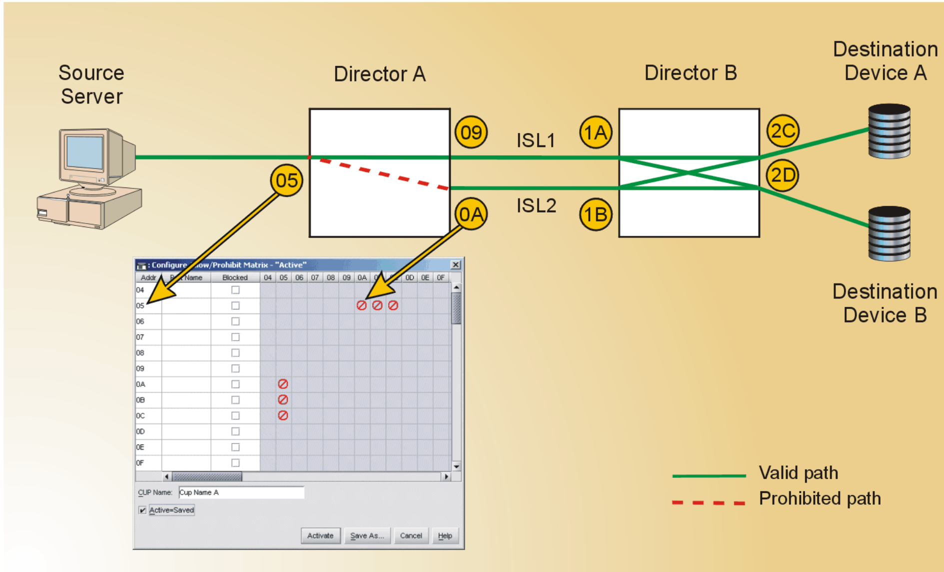

Figure 58: PDCM Array - Example Problem

A PDCM array configured for Director **A** prohibits logical port address **05** from communicating with logical port addresses **0A**, **0B**, and **0C**. No PDCM array is configured for Director **B**. The PDCM array configured for Director **A** prohibits the source server from transmitting or receiving data across ISL **2**. However, internal route tables on both directors indicate a valid server-to-destination device path across ISL **1**.

A problem arises when the source server transmits Class 3 Fibre Channel data to devices across ISL **1**, consuming the ISL bandwidth. Destination devices are unaware of the PDCM array configured at Director **A** and transmit frames back to the server across ISL **2**. Because the server is prohibited from communicating across this ISL, Class 3 Fibre Channel frames are discarded without generating a busy (BSY) frame, reject (RJT) frame, or otherwise notifying the destination devices. The server receives no response from destination devices and times out. Thus, a server or device failure is indicated when in fact the problem is a user-defined prohibited connection.

Preferred Path

The preferred path option allows a user to specify and configure one or more ISL data paths between multiple directors or switches in a fabric. At each fabric element, a preferred path consists of a source port on the director or switch being

configured, an exit port on the director or switch, and the domain ID of the destination director or switch. Each participating director or switch must be configured as part of a desired path. The following rules apply when configuring a preferred path:

- The switch domain ID must be set to **Insistent**.
- Domain IDs range between **1** through **31**.
- Source and exit port numbers are limited to the range of ports available on the director or switch.

For each source port, only one path is defined to each destination domain ID.

As an example, refer to the three-director preferred path illustrated in Figure 59. A preferred path is configured between a source server and destination device (**A** or **B**), traversing Director **1**, Director **2**, and Director **3**. To configure the preferred path through the first director:

1. Select the **Preferred Path** option from the *Element Manager* application's **Configure** menu. The **Configure Preferred Paths** dialog box displays.
2. Click **Add**. The **Add Preferred Path** dialog box displays (bottom of Figure 59).
3. For the director entry port, type **14** in the **Source Port** field. For the director exit port, type **45** in the **Exit Port** field. For the destination device (Director **3**), type **22** in the **Destination Domain ID** field.
4. Click **OK** to save the path configuration and close the dialog box.

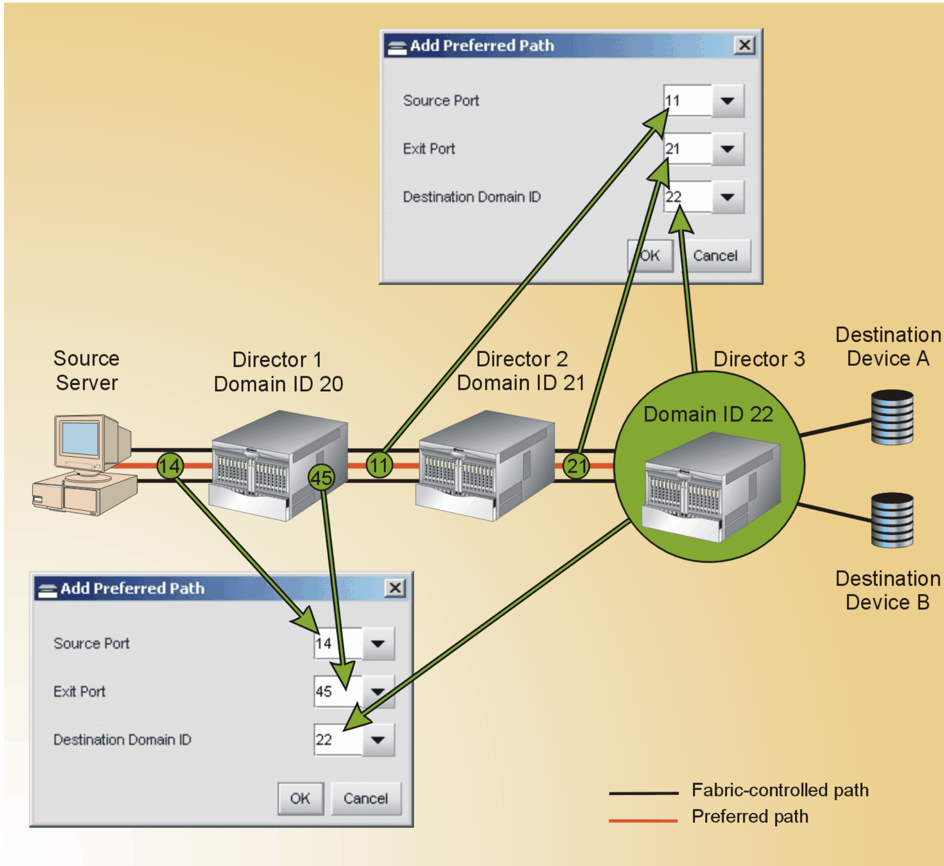

Figure 59: Preferred Path Configuration

This procedure only specifies that data will enter and exit Director **1** through specific ports on the path to Director **3**. The procedure must be repeated at the second director as follows:

1. Select the **Preferred Path** option from the *Element Manager* application's **Configure** menu. The **Configure Preferred Paths** dialog box displays.
2. Click **Add**. The **Add Preferred Path** dialog box displays (top of Figure 59).
3. For the director entry port, type **11** in the **Source Port** field. For the director exit port, type **21** in the **Exit Port** field. For the destination device (Director **3**), type **22** in the **Destination Domain ID** field.

4. Click **OK** to save the path configuration and close the dialog box.

Activating a preferred path can result in receipt of out-of-order frames (especially in FICON environments) if the path differs from the current path, if input and output (I/O) are active from the source port, and if congestion is present on the current path.

To avoid problems in FICON environments, vary associated channel path identifiers (CHPIDs) temporarily offline, configure the preferred path, and vary the CHPIDs back online.

Zoning

Directors and switches support a user configuration that partitions attached devices into restricted-access groups called zones. Devices in the same zone can recognize and communicate with each other through switched port-to-port connections. Devices in separate zones cannot recognize name server or route table information and therefore cannot communicate with each other. Figure 60 illustrates a Director 2/64 with three zones (four devices per zone).

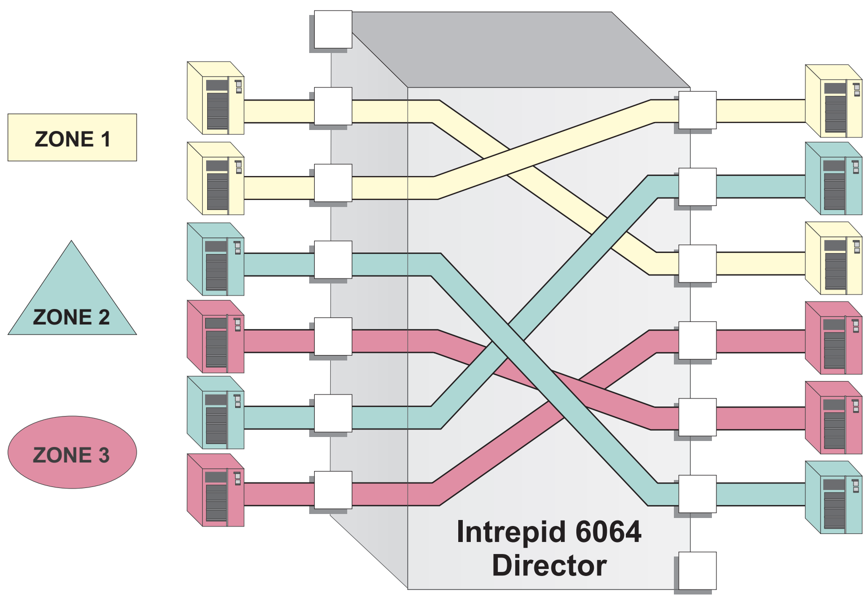

Figure 60: Product Zoning

Zoning is enabled and enforced by one of the following processes:

- **Software-enforced zoning** — For earlier versions of director or switch firmware (prior to version 6.0), the device configuration on a fabric element enforces zoning by limiting access to name server information in response to a device query. Only devices in the same zone as the requesting device are returned in the query response. This type of zoning is also called name server zoning or soft zoning.
- **Hardware-enforced zoning** — For later versions of director or switch firmware (Version 6.0 and later), the device configuration on a fabric element enforces zoning by programming route tables that strictly prevent Fibre Channel traffic between devices that are not in the same zone. This type of zoning is also called hard zoning.

Zones are configured through the *HAFM* application by authorizing or restricting access to name server or route table information (depending on the firmware release level) associated with device N_Ports that attach to director or switch F_Ports.

Benefits of Zoning

System administrators create zones to increase network security measures, differentiate between operating systems, and prevent data loss or corruption by controlling access between devices (such as servers and data storage units) or between separate user groups (such as engineering or human resources). Zoning allows an administrator to establish:

- Logical subsets of closed user groups. Administrators can authorize access rights to specific zones for specific user groups, thereby protecting confidential data from unauthorized access.
- Barriers between devices that use different operating systems. For example, it is often critical to separate servers and storage devices with different operating systems because accidental transfer of information from one to another can delete or corrupt data. Zoning prevents this by grouping devices that use the same operating systems into zones.
- Groups of devices that are separate from devices in the rest of a fabric. Zoning allows certain processes (such as maintenance or testing) to be performed on devices in one group without interrupting devices in other groups.
- Temporary access between devices for specific purposes. Administrators can remove zoning restrictions temporarily (for example, to perform nightly data backup), then restore zoning restrictions to perform normal processes.

Configuring Zones

Zoning is configured through the *HAFM* application by authorizing or restricting access to name server information associated with device node ports (N_Ports) that attach to director or switch fabric ports (F_Ports). A device N_Port can belong to multiple zones. Zoning is configured by:

- The eight-byte (64-digit) World Wide Name (WWN) assigned to the HBA or Fibre Channel interface installed in the device connected to the director or switch (recommended method).

> **Caution:** If zoning is implemented by WWN, removal and replacement of a device HBA or Fibre Channel interface (thereby changing the device WWN) disrupts zone operation and may incorrectly exclude a device from a zone.

- The domain identification (ID) and physical port number of the director or switch port to which the device is attached.

> **Caution:** If zoning is implemented by port number, a change to the director or switch fiber-optic cable configuration disrupts zone operation and may incorrectly include or exclude a device from a zone.

A zone contains a set of attached devices that can access each other. Zones are grouped into *zone sets*. A zone set is a group of zones that is enabled (activated) or disabled across all directors and switches in a multi-switch fabric. Only one zone set can be enabled at one time. Zone members are defined and zones or zone sets are created using the *HAFM* application. HP products support the following zoning features:

- **Zone members** — The maximum number of members configurable for a zone is 4,096.
- **Number of zones** — The maximum number of configurable zones in a zone set is 1,023 (1,024 including the default zone).
- **Number of zone sets** — The maximum number of configurable zones sets is 64.
- **Active zone set** — The zone set that is active across all directors and switches in a multi-switch fabric. For the active zone set:

— When a specific zone set is activated, that zone set replaces the active zone set.

— If the active zone set is disabled, all devices attached to the fabric become members of the default zone.

— All devices not included as members of the active zone set are included in the default zone.

- **Default zone** — The default zone consists of all devices not configured as members of a zone in the active zone set. If there is no active zone set, all devices attached to the fabric are in the default zone. For the default zone:

  — The default zone is enabled or disabled separately from the active zone set.

  — If the default zone is enabled, all devices not in a specified zone are included in the default zone and can communicate with each other.

  — If the default zone is disabled and there is no active zone set, the zoning feature is completely disabled for the fabric and no devices can communicate with each other.

  — All devices are considered to be in the default zone if there is no active zone set.

- **RSCN service requests** — Registered state change notification (RSCN) service requests are transmitted to all N_Ports attached to the director or switch when the zoning configuration is changed.

- **Broadcast frames** — Class 3 broadcast frames are transmitted to all N_Ports attached to the director or switch, regardless of zone membership.

Joining Zoned Fabrics

Directors and edge switches are linked through ISLs to form multi-switch fabrics. In a multi-switch fabric, the active zoning configuration applies to the entire fabric. Any change to the configuration applies to all directors and switches in the fabric.

When fabrics attempt to join, participating fabric elements exchange active zone configurations and determine if their configurations are compatible. If the configurations are compatible, the fabrics join. The resulting configuration is a single zone set containing zone definitions from each fabric.

If the configurations cannot merge, E_Ports that form the ISL for each fabric element become segmented. The ports cannot transmit data frames between attached switches (Class 2 or 3 traffic) but can transmit control frames (Class F traffic).

Zoning configurations are compatible if there are no duplicate domain IDs, the active zone set name is the same for each fabric (or switch in the fabric), and zones with the same names in each fabric have identical members.

Factors to Consider When Implementing Zoning

Consider the following factors when planning to implement zoning for one or more directors or switches in the enterprise. In particular, consider the implications of zoning within a multi-switch fabric.

- **Reasons for zone implementation** — Determine if zoning is to be implemented for the enterprise. If so, evaluate if the purpose of zoning is to differentiate between operating systems, data sets, user groups, devices, processes, or some combination thereof. Plan the use of zone members, zones, and zone sets accordingly.
- **Zone members specified by port number or WWN** — Determine if zoning is to be implemented by port number or WWN. Because changes to port connections or fiber-optic cable configurations disrupt zone operation and may incorrectly include or exclude a device from a zone, zoning by WWN is recommended. However, if zoning is implemented by WWN, removal and replacement of a device’s HBA or Fibre Channel interface disrupts zone operation and will exclude a new device from a zone unless the device is added to the zone set.
- **Zoning implications for a multi-switch fabric** — For a multi-switch fabric, zoning is configured on a fabric-wide basis, and any change to the zoning configuration is applied to all switches in the fabric. To ensure zoning is consistent across a fabric, there can be no duplicate domain IDs, the active zone set name must be consistent, and zones with the same name must have identical elements. Ensure these rules are enforced when planning zones and zone sets, and carefully coordinate the zoning and multi-switch fabric tasks.

Obtaining Professional Services

Planning and implementing a multi-switch fabric can be a complex and difficult task. HP recommends that you obtain planning assistance from our professional services organization before implementing a fabric topology.

Server and Storage-Level Access Control

To enhance the access barriers and network security provided by zoning through the director or switch, security measures for SANs should also be implemented at servers and storage devices.

Server-level access control is called *persistent binding*. Persistent binding uses configuration information stored on the server and is implemented through the server's HBA driver. The process binds a server device name to a specific Fibre Channel storage volume or logical unit number (LUN) through a specific HBA and storage port WWN.

For persistent binding:

- Each server HBA is explicitly bound to a storage volume or LUN and access is explicitly authorized (access is blocked by default).
- The process is compatible with OSI standards. The following are transparently supported:
  - Different operating systems and applications.
  - Different storage volume managers and file systems.
  - Different fabric devices, including disk drives, tape drives, and tape libraries.
- If the server is rebooted, the server-to-storage connection is automatically re-established.
- The connection is bound to a storage port WWN. If the fiber-optic cable is disconnected from the storage port, the server-to-storage connection is automatically re-established when the port cable is reconnected. The connection is also automatically re-established if the storage port is cabled through a different director or switch port.

Access control can also be implemented at the storage device as an addition or enhancement to redundant array of independent disks (RAID) controller software. Data access is controlled within the storage device, and server HBA access to each LUN is explicitly limited (access is blocked by default).

Storage-level access control:

- Provides control at the storage port and LUN level and does not require configuration at the server.
- Supports a heterogeneous server environment and multiple server paths to the storage device.

- Is typically proprietary and protects only a specific vendor's storage devices. Storage-level access control may not be available for many legacy devices.

Security Best Practices

When implementing a enterprise data security policy, establish a set of best practice conventions using methods described in this section in the following order of precedence (most restrictive listed first):

1. **SANtegrity Binding** — The SANtegrity Binding feature is recommended for large and complex SANs with fabrics and devices provided by multiple OEMs or that intermix FCP and FICON protocols. The feature is required for FICON-cascaded high-integrity SANs. SANtegrity Binding includes:
   - Fabric binding (configured and enabled through the *HAFM* application) that allows only user-specified directors or switches to attach to specified fabrics in a SAN.
   - Switch binding (configured and enabled through the *Element Manager* application) that allows only user-specified devices and fabric elements to connect to specified director or fabric switch ports.

   SANtegrity Binding explicitly prohibits connections that are not user configured (unauthorized ISLs or device connections *do not* initialize and devices *do not* log in) and takes precedence over allowed connectivity in PDCM arrays, allowed connectivity through hard or soft zoning, preferred path configurations, or device-level access control.
2. **PDCM arrays** — In FICON environments, connectivity control is configured and managed at the director or switch level using a PDCM array, where a user specifies which logical port addresses are allowed or prohibited from connecting with each other, including E_Port connectivity.

   Port-to-port connectivity is hardware enforced at each fabric element, and explicitly prohibited connections take precedence over allowed connectivity through hard or soft zoning, preferred path configurations, or device-level access control. However, a connection allowed through a PDCM array may be prohibited through SANtegrity Binding.
3. **Hardware-enforced zoning** — The function of hard zoning is to ensure that route tables are programmed at each fabric element that explicitly allow devices to communicate *only* if the devices are in the same zone. Zoning configurations are hardware-enforced at each fabric element source port. Hard zoning impacts devices only and does not prohibit E_Port (ISL) connectivity.

Devices in common zones can be prohibited from communicating through SANtegrity Binding or PDCM arrays, but hard zoning takes precedence over preferred path configurations, allowed connectivity through soft zoning, or device-level access control.

4. **Preferred path** — A preferred path provides soft control of fabric routing decisions on a switch-by-switch or port-by-port basis. The path instructs a fabric to use a preferred exit port out of a director or switch for a specified receive port and target domain.

   If a preferred path is prohibited by SANtegrity Binding, PDCM arrays, or hard zoning, the path is not programmed. In addition, if a preferred path is not a shortest path as calculated by Dijkstra's fibre shortest path first (FSPF) algorithm, the preferred path is not programmed. However, preferred paths do take precedence over dynamic load balancing enabled through the OpenTrunking feature, soft zoning, or device-level access control.

   In general, preferred paths should be configured to influence predictable or well-known Fibre Channel traffic patterns for load balancing or distance extension applications.

5. **Software-enforced zoning** — When a device queries the name server of a fabric element for a list of other attached devices, soft zoning ensures only that a list of devices is in the same zone as the requesting device it returned. Soft zoning only informs a device about authorized zoning configurations; it does not explicitly prohibit an unauthorized connection. Connectivity configured through SANtegrity Binding, PDCM arrays, hardware-enforced zoning, and preferred paths takes precedence over soft zoning.

6. **Device-level access control** — Persistent binding and storage access control can be implemented at the device level as an addition or enhancement to other security features (SANtegrity Binding, PDCM arrays, zoning, and preferred paths) that are more explicitly enforced.

Security methods described in this section work in parallel with each other and are allowed to be simultaneously enabled and activated. Users are responsible for security configuration and operation within the constraints and interactions imposed by their fabric design and the methods described here.

Because incompatible security configurations can cause unintended connectivity problems or shut down Fibre Channel traffic in a fabric, it is imperative that users study and understand the interactions between SANtegrity Binding, PDCM arrays, zoning, preferred paths, and device-level access control. It is recommended to follow the best practices listed here in order of precedence.

Logically work in sequence from the most restrictive method to the least restrictive method, ensuring the most restrictive routing or connectivity paths override all other paths.

Optional Features

HP offers several operating features that are available for the switch as customer-specified options. Available features include:

- **Open Systems Management Server** or **FICON Management Server** — Inband director or switch management is provided through purchase of the OSMS or FMS feature.

Note: The Edge Switch 2/24 does not support out-of-band management through FMS.

- **Flexport Technology** — The Flexport technology feature is a hardware port expansion kit that allows customers to upgrade switch capacity on demand in eight-port increments. A Flexport technology switch is delivered at a discount without all the ports enabled. When additional port capacity is required, the remaining ports are enabled (in eight-port increments) through purchase of this feature.

Note: The Director 2/64 and Director 2/140 do not support the Flexport technology feature.

- **SANtegrity Binding** — Purchase and enabling of this feature enhances security in SANs that contain a large and mixed group of fabrics and attached devices.
- **Open Trunking** — Purchase and enabling of this feature provides dynamic load balancing of Fibre Channel traffic across multiple ISLs.
- **Full volatility** — Purchase and enabling of this feature ensures that no Fibre Channel frames are stored after a director or switch is powered off or fails, and a memory dump file (that possibly includes classified frames) is not included as part of the data collection procedure.
- **CNT WAN support** — This feature is included *only* in software maintenance release 4.02.00 and is required to allow an Edge Switch 2/12 or Edge Switch 2/24 to communicate with Computer Network Technologies (CNT) UltraNet Edge storage routers.

Note: Director 2/64, Director 2/140, Edge Switch 2/16, and Edge Switch 2/32 do not provide CNT wide area network (WAN) support.

- **Element Manager application —** This feature enables director or switch management through the Element Manager user interface. Directors and switches are delivered with the application enabled for a 31-day grace period. Before grace period expiration, the application must be reactivated through a PFE key.

After purchasing a feature, obtain the required feature enablement key through your HP marketing representative. A feature key is an alphanumeric string consisting of both uppercase and lowercase characters. The total number of characters may vary. The feature key is case sensitive and must be entered exactly, including dashes. The following is an example of a feature key format:

XxXx-XXxX-xxXX-xX.

Inband Management Console Access

Inband management console access (through a Fibre Channel port) is provided by enabling user-specified features that allow Open Systems (OSMS) or FICON (FMS) host control of a director or switch. The features are mutually exclusive; only one can be installed at a time.

Open Systems Management Server

When the OSMS feature key is enabled with the *Element Manager* application, host control and management of the director or switch is provided through an open systems interconnection (OSI) device attached to a product port.

When implementing inband product management through an OSI connection, plan for the following minimum host requirements:

- Connectivity to an OSI server with a product-compatible host bus adapter (HBA) that communicates through the Fibre Channel common transport (FC-CT) protocol.
- Installation of a storage network management application on the OSI server. Management applications include Veritas SANPoint Control (version 1.0 or later) or Tivoli NetView (version 6.0 or later).

For information about product-compatible HBAs, third-party SAN management applications, and minimum OSI server specifications, refer to the HP web site at http://www.hp.com.

FICON Management Server

When the FMS feature key is enabled with the *Element Manager* application, host control and management of the director or switch is provided through a server attached to a product port. The server communicates with the product through a FICON channel.

When implementing inband product management through a FICON channel, plan for the following minimum host requirements:

- Connectivity to an IBM S/390 Parallel Enterprise Server (Generation 5 or Generation 6), with one or more FICON channel adapter cards installed, using System Automation for Operating System/390 (SA OS/390) for native FICON, version 1.3 or later, plus service listed in the appropriate preventive service planning (PSP) bucket. The PSP bucket upgrade is HKYSA30.

  The minimum OS/390 level for a director or switch without the control unit port (CUP) feature is version 2.6, plus service listed in PSP bucket upgrade 2032, device subset 2032OS390G5+. The minimum OS/390 level for a director or switch with the CUP feature is version 2.1, plus service listed in the preceding PSP bucket for that function.
- Connectivity to an IBM eServer zSeries 800 (z800), zSeries 900 (z900), or zSeries 990 (z990) processor, with one or more FICON or FICON Express channel adapter cards installed, using the z/OS operating system, version 1.1 or later.
- A host-attached Hardware Management Console. The console runs the Hardware Management Console application (HWMCA) and is the operations and management PC platform for S/390 or zSeries servers.

For additional information, refer to the IBM publication, *System Automation for OS/390, Operations* (GC28-1550).

Flexport Technology

The Edge Switch 2/12, Edge Switch 2/32, and Edge Switch 2/24 can be purchased at a discount without all Fibre Channel ports enabled. The Flexport technology feature is a hardware port expansion kit that allows customers to upgrade switch capacity on demand in eight-port increments.

Flexport technology kits are available to upgrade the following:

- Edge Switch 2/12 from 4 to 8 ports and/or 8 to 12 ports.
- Edge Switch 2/24 from 8 to 16 ports and/or from 16 to 24 ports.
- Edge Switch 2/32 from 16 to 24 ports and/or 24 to 32 ports.

Each port expansion kit includes eight SFP optical transceivers, upgrade instructions, and a feature key that enables the added port capacity through the *Element Manager* application.

SANtegrity Binding

SANtegrity Binding is a feature that enhances data security in large and complex SANs that have numerous fabrics and devices provided by multiple original equipment manufacturers (OEMs). This feature allows or prohibits director or switch attachment to fabrics (fabric binding) and Fibre Channel device attachment to directors or switches (switch binding). The SANtegrity binding feature includes:

- **Fabric binding** — Using the fabric binding feature, an administrator allows only specified directors or switches to attach to specified fabrics in a SAN. This provides security from accidental fabric merges or potential fabric disruption when multiple fabrics segment because they cannot merge. This feature is managed through the *HAFM Manager* application.
- **Switch binding** — Using the switch binding feature, an administrator allows only specified devices and fabric elements to connect to specified director or fabric switch ports. This provides security in environments that include a large number of devices by ensuring that only the intended set of devices attaches to a director or switch. This feature is managed through the *Element Manager* application.

Enterprise Fabric Mode

Although *Enterprise Fabric Mode* is not a keyed feature, it is integral to SANtegrity Binding operation. *Enterprise Fabric Mode* must be enabled through the *HAFM Manager* application before fabric binding and switch binding can operate. Enterprise Fabric Mode also enables the following parameters:

- **Rerouting delay** — If a fabric topology changes, directors or switches calculate a new least-cost data transfer path through a fabric, and routing tables immediately implement that path. This may result in Fibre Channel frames being delivered to a destination device out of order, because frames transmitted over the new (shorter) path may arrive ahead of previously transmitted frames that traverse the old (longer) path. When enabled, the rerouting delay parameter ensures that frames are delivered through a fabric in the correct order.

- **Domain RSCNs** — Domain RSCNs provide connectivity information to all HBAs and storage devices attached to a fabric. RSCNs are transmitted to all registered device N_Ports attached to a fabric if either a fabric-wide event or zoning configuration change occurs.
- **Insistent domain ID** — When this parameter is enabled, the domain ID configured as the preferred domain ID for a director or switch becomes the active domain ID when the fabric initializes. A static and unique active domain identification is required by the fabric binding feature because the feature's fabric membership list identifies fabric elements by WWN and domain ID. If a duplicate preferred domain ID is used, then insisted upon, a warning occurs and the affected director or switch cannot be added to the membership list.

SANtegrity Binding Planning Considerations

Fabric and switch binding enhance data security by controlling and monitoring director, fabric switch, and device connectivity. The name server zoning feature also provides data security by partitioning devices into restricted-access zones. Use of the SANtegrity Binding and zoning features in conjunction with each other must be carefully planned and coordinated. Refer to “Zoning” on page 154 for additional information about zoning.

It is recommended you obtain planning assistance from the HP professional services organization before implementing the SANtegrity Binding feature with director or switch zoning, especially for multiple fabrics.

Open Trunking

Open Trunking is a feature that optimizes ISL bandwidth use in a fabric environment. This feature monitors Fibre Channel data rates (congestion and BB_Credit starvation) through multiple ISLs, dynamically applies a Dijkstra FSPF networking algorithm to calculate the optimum path between fabric elements, and load balances Fibre Channel traffic (from congested links to uncongested links) accordingly. Open Trunking is shown in Figure 61.

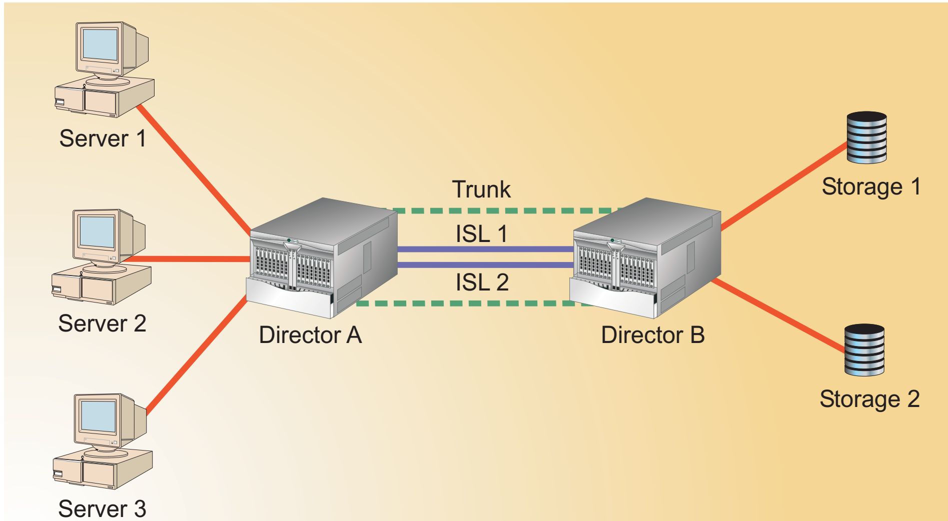

Figure 61: Open Trunking configuration

The figure illustrates two Director 2/64 directors connected by two ISLs. Three servers use the ISLs to communicate with two storage devices. Without trunking, servers **1** through **3** route Fibre Channel traffic to director B without regard to any data rates. A possible scenario is that servers **1** and **2** route high data rate traffic through ISL **1** to storage device **1** (ISL oversubscription) and server **3** routes low data rate traffic through ISL **2** to storage device **2** (ISL undersubscription).

Note that preferred path configurations are more restrictive than and take precedence over Open Trunking. Even if Open Trunking is enabled, no attempt is made to reroute traffic away from a preferred path, even if the path is congested or BB_Credit starved.

Full Volatility

Full volatility is a feature (available on directors and switches with firmware version 6.0 and later) that supports military, classified, or other high-security environments that require that Fibre Channel data not be retained by the director or switch after power-off or failure.

When a director or switch (without the full-volatility feature installed) powers off or fails, a dump file is written to non-volatile random-access memory (NV-RAM). This dump file retains the last 30 Fibre Channel frames transmitted from the embedded port and the last four frames transmitted to the embedded port.

These Fibre Channel frames are then written to CD and included as part of the data collection procedure. This process constitutes a security breach if the frame data includes classified information.

With the full-volatility feature installed and enabled, no frame data is stored, and the NV-RAM dump does not occur when the director or switch powers off or fails. Although this feature limits the diagnostic information available for fault isolation and resolution, the majority of failures are resolved without the dump file.

CNT WAN Support

Fibre Channel-based SANs are typically implemented as isolated networks accessible only from local servers connected through a Fibre Channel fabric. Connectivity between devices is usually limited to about ten kilometers or fewer. Many companies are striving to interconnect isolated SANs and consolidate computer resources through WAN extension technology. Therefore, edge switches deployed as part of a core-to-edge fabric often require WAN connectivity.

This connectivity is provided by the CNT WAN support feature, included only in software maintenance release 4.02.00. With this feature installed and enabled, an Edge Switch 2/12 or Edge Switch 2/24 can communicate with CNT UltraNet Edge storage routers (WAN gateways).

Element Manager Application

The Element Manager feature allows director or switch management through an *Element Manager* application GUI. A director or switch is delivered with the application enabled for a 31-day grace period. Before grace period expiration, the application must be reactivated and enabled through a PFE key.

During the grace period, a **No Feature Key** dialog box (Figure 62) displays when the *Element Manager* application is accessed. Click **OK** to close the dialog box and open the application.

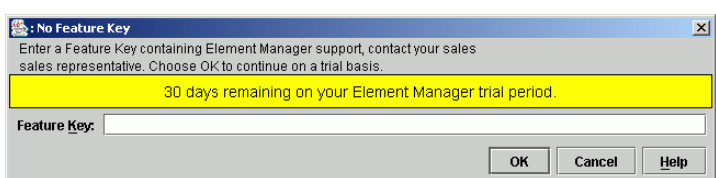

Figure 62: No Feature Key Dialog Box

Note: The Edge Switch 2/12 is not supported by the HAFM appliance.

In addition, the message **Element Manager license key has not been installed - Please follow up instructions to update permanent key** is splashed across views, indicating the Element Manager PFE key must be installed. The **Hardware View** (Figure 63) for an Edge Switch 2/24 is shown as an example.

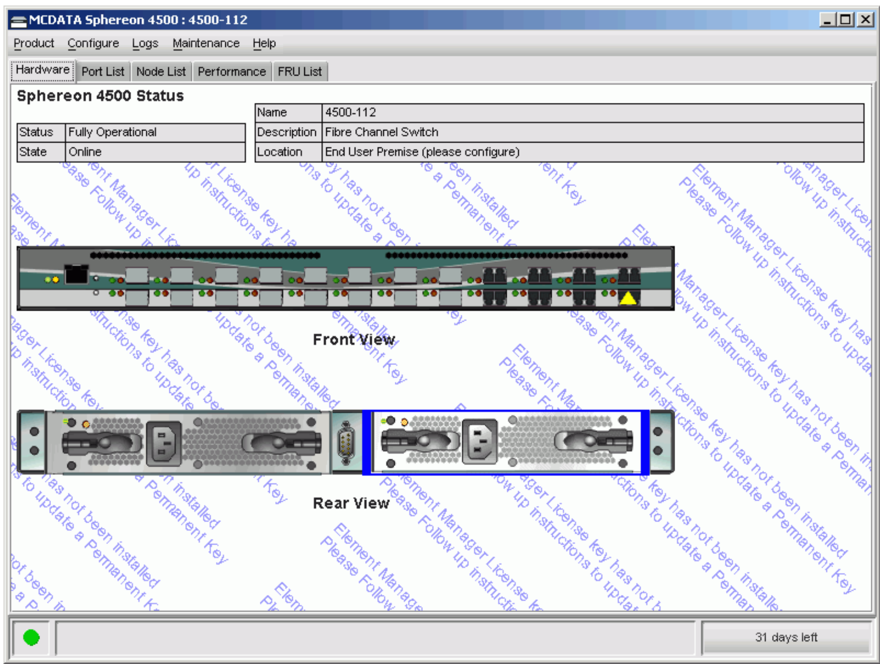

Figure 63: Hardware View (with Element Manager Message)

Configuration Planning Tasks

5

This chapter describes configuration planning tasks to be performed before installing the High Availability Fabric Manager (HAFM) server and one or more Director 2/64s, Director 2/140s, Edge Switch 2/12s, Edge Switch 2/16s, Edge Switch 2/24s, or Edge Switch 2/32s in a storage area network (SAN) configuration.

Note: The Edge Switch 2/12 is not supported by the HAFM appliance.

The following planning tasks are described in this chapter.

Task 1: Prepare a Site Plan

For each director, switch, or equipment rack installed, design a site plan that provides efficient work flow, operator convenience and safety, and adequate service clearances for the equipment rack.

A customer manager should review the site plan with a service representative and consider:

- Location and relationship of the physical facilities, such as walls, doors, windows, partitions, furniture, and telephones.
- Proximity of the director or switch to servers and storage peripherals, and if a multi-switch fabric is to be enabled, proximity of participating fabric elements to each other.
- Location of at least one analog phone line to aid in installation and serviceability.
- Availability of Ethernet local area network (LAN) connections and cabling to support remote user workstation and simple network management protocol (SNMP) management station access. Remote user and SNMP workstations are optional.
- Equipment rack locations, Ethernet cabling, and the Internet Protocol (IP) addressing scheme to support optional rack interconnection and HAFM appliance consolidation.
- Power requirements, including an optional uninterruptable power supply (UPS).
- Lengths of power cables and location of electrical outlets (for directors, switches, and the HAFM appliance) having the proper phase, voltage, amperage, and ground connection.

WARNING: An insulated grounding conductor identical in size, insulating material, and thickness to the grounded and ungrounded branch-circuit supply conductors (except it is green, with or without one or more yellow stripes) shall be installed as part of the branch circuit supplying the product. The grounding conductor described shall be connected to ground at the product, or if supplied by a separately derived system, at the supply transformer or motor generator. The plug receptacles near the product shall all be a grounding type, and grounding conductors serving these receptacles shall be connected to ground at the product.

- Security necessary to protect the installation's physical integrity, while maintaining accessibility to the director or switch.
- Equipment rack front and rear service clearances, operator clearances, and maintenance access clearances.
- Weight of an equipment rack. Either multiple persons or a lift must be available during installation to remove the rack from the packing crate.
- Heat dissipation, temperature, and humidity requirements.

Complete the planning checklists under this task. The checklists provide detailed planning activities and provide space for a planned completion date for each activity. The customer's management information system (MIS) project manager should examine the checklists and determine the personnel and resources required for completing planning and installation tasks.

Customer personnel might be used from the following functional areas:

- Systems programming personnel to update input/output (I/O) definitions to identify directors and switches.
- Ethernet management personnel to obtain IP addresses, gateway addresses, and subnet masks for directors, switches, and the HAFM appliance and a domain name system (DNS) host name for the HAFM appliance.
- Facilities planning personnel to outline the facility floor plan and to arrange for electrical wiring, receptacles, and telephone lines.
- Installation planning personnel to determine fiber-optic and Ethernet cabling requirements and routing requirements and to plan connectivity between each director, switch, and attached device.
- Trainers to determine training and education needs for operations, administration, and maintenance personnel.
- Administrators to determine director port names and WWN nicknames, identify attached devices, and assign password levels and user names for director and switch access.

Table 4 lists physical planning and hardware installation tasks and includes the activity, task owner, due date, and comments.

Table 4: Physical Planning and Hardware Installation Tasks

Activity	Task Owner	Due Date	Comments
Locate the physical facilities.			
Connect the facility alternating current (AC) power circuits.			If more than one director or switch, consider separate power circuits for availability.
Obtain an uninterruptable power supply (optional).			Recommended.
Obtain an outside-access phone line.			Telephone for support personnel.
Order the equipment rack with one or more HP products.			
Order Fibre Channel devices and peripherals.			
Install Fibre Channel devices and peripherals.			
Route fiber-optic jumper cables.			
Determine proximity of the equipment rack (with directors and switches) to attached devices (multimode shortwave laser or single-mode longwave laser).			250 meters (2.125 Gbps) for 50/125 mm multimode cable. 150 meters (2.125 Gbps) for 62.5/125 mm multimode cable. 10, 20, or 35 kilometers for 9/125 mm single-mode cable.
Order and deliver fiber-optic cables.			Cables are purchased by the customer separately. Plan to have them arrive and laid out before equipment rack delivery.
Set up local area network (LAN) connections for directors, switches, and the HAFM appliance.			
Set up LAN connections to corporate intranet for remote workstation access (optional).			

Table 5 lists operational setup tasks and includes the task owner, due date, and comments.

Table 5: Operational Setup Tasks

Activity	Task Owner	Due Date	Comments
Obtain IP address and subnet mask.			HAFM appliance (if installing on a LAN with non-HP devices). Directors and switches (if installing on a LAN with non-HP devices). Remote user workstation (optional). Simple network management protocol (SNMP) management stations (optional).
Obtain gateway addresses for router or other gateway devices on company LAN.			To configure on HAFM appliance and products (if installing on a LAN with non-HP devices).
Assign host names.			HAFM appliance and products (optional).
Add host name to DNS database.			HAFM appliance and products.
Determine what level of *HAFM* application user rights are to be used for up to 16 users.			
Determine if inband management of the director or switch is to be used and the type (FICON or Open Systems).			HAFM appliance and Fibre-Channel-attached server peripheral (optional).
Determine if the call-home feature is to be used.			
Determine if the e-mail notification feature is to be used.			Obtain e-mail addresses for event notification and identify the e-mail server.

Table 5: Operational Setup Tasks (Continued)

Activity	Task Owner	Due Date	Comments
Determine SNMP access to directors and switches.			Obtain SNMP trap recipient IP addresses. Determine SNMP information required (generic and product-specific). Determine if write permission is required for modifying SNMP variables.
Determine if a multi-switch fabric is to be implemented.			
Determine if the zone management feature is to be used.			
Introduce staff to *HAFM* and *Element Manager* applications.			
Introduce staff to remote session parameters.			
Introduce staff to product recovery concepts and messages.			
Assign port names.			
Configure extended distance (10 to 100 km) ports.			
Configure link incident alerts.			
Configure Ethernet events.			

Task 2: Plan Fibre Channel Cable Routing

Plan for sufficient single-mode fiber-optic and multimode fiber-optic cabling to meet the connectivity requirements for all Fibre Channel servers and devices. If a multi-switch fabric is to be enabled, plan for sufficient fiber-optic cabling to meet interswitch link (ISL) connectivity requirements.

Plan for at least one meter (39.37 inches) of fiber-optic cable inside the equipment rack for routing to product Fibre Channel ports as required. Plan for an additional 1.5 meters (5 feet) of cable outside the rack to provide slack for service clearance, limited rack movement, and inadvertent cable pulls.

WARNING: Director and switch non-open fiber control (non-OFC) laser transceivers are designed and certified for use only with fiber-optic cable and connectors with characteristics specified by HP. Use of other connectors or optical fiber can result in emission of laser power levels capable of producing injury to the eye if viewed directly. Use of non-specified connectors or optical fiber can violate the Class 1 laser classification.

In addition, consider the following when planning cable routing:

- The need for additional fiber-optic cables could grow rapidly. Consider installing cable with extra fibers, especially in hard-to-reach places like underground trenches. Consider locating the equipment rack near a fiber-optic patch panel.
- Follow proper procedures when moving an installed equipment rack to prevent cable or connector damage.

Task 3: Consider Interoperability with Fabric Elements and End Devices

HP conducts a substantial level of testing to ensure director and switch interoperability with fabric elements and end devices provided by multiple original equipment manufacturers (OEMs). New devices are tested and qualified on a continual basis. Contact your HP representative for the latest information about fabric element, server, host bus adapter (HBA), and device interoperability.

Consider whether to set the director or switch to Open Systems management style or FICON management style. This setting affects only the management style used to manage the product; it does not affect port operation. Open-systems interconnection (OSI) devices can communicate with each other if the product is set to FICON management style, and Fibre Connection (FICON) devices can communicate with each other if the product is set to Open Systems management style.

Be aware that:

- When a director or switch is set to Open Systems management style, a traditional Fibre Channel fabric consisting of multiple domains (fabric elements) is supported. Inband management through the open-systems management server (OSMS) is also supported.
- When a director or switch is set to FICON management style, only a single domain (fabric element) is supported. Inband management through the FICON management server (FMS) is also supported. When operating in FICON management style, ports are set to F_Port operation, thus eliminating E_Port, ISL, and multiswitch fabric capabilities.

Note: If the FICON management server feature is enabled, the default management style is FICON. Open Systems management style cannot be enabled.

Task 4: Plan Console Management Support

Plan to implement one or more of the following methods to provide console management and support for directors and switches:

- **HAFM appliance** — The rack-mounted HAFM appliance is used for product installation, initial software configuration, changing the configuration, and monitoring product operation.
  - When the *HAFM* application and *Element Manager* applications are installed on the HAFM appliance, the server is used as a local user workstation.
  - The HAFM appliance can support up to 48 managed HP products.
  - Managed directors and switches can be powered off and on without the HAFM appliance.
  - An HAFM appliance failure does not cause an operating director or switch to fail.
  - The HAFM appliance is fully operational, even if there is no user logged in to the Windows 2000 operating system. The HAFM appliance allows remote users to log in and continues to monitor products in the background.
- **Remote user workstations** — If remote access to the HAFM appliance is required, plan to install user workstations with the *HAFM* and *Element Manager* applications configured. Administrators can use these remote workstations to configure and monitor directors and switches. Up to 25 HAFM sessions can be simultaneously active (one local from the HAFM appliance and 24 remote). Sessions from remote user workstations are disabled if the HAFM appliance is powered off.
- **Inband management support** — If inband console management of a director or switch is required, plan for a Fibre Channel port connection that communicates with the attached server.

  If director or switch management through an OSI server is planned, ensure that the OSMS feature key is ordered with the *Element Manager* application. This feature enables host control of the product from an OSI server attached to a Fibre Channel port. Ensure that the server meets minimum specifications and that a product-compatible HBA and appropriate operating system or SAN management application is available.

If director or switch management through an IBM host is planned, ensure that the FMS feature key is ordered with the *Element Manager* application. This feature key enables host control of the product from an IBM System/390 or zSeries 900 Parallel Enterprise server attached to a Fibre Channel port.

- **Web server interface** — If Internet access to a director or switch Embedded Web Server interface is required, plan for access to an analog phone line. Internet access to the Web server interface is not provided by the HAFM appliance.

Task 5: Plan Ethernet Access

Directors and the HAFM appliance can be ordered in an HP-supplied equipment rack in accordance with customer specifications; however, you may need to:

- **Connect equipment racks** — Customer-supplied Ethernet hubs in multiple equipment racks can be connected to provide HAFM appliance access to up to 48 managed HP products. Racks can be placed at any distance up to the limit of the 10/100 megabit per second (Mbps) LAN segment.
- **Consolidate HAFM appliance operation** — If HAFM appliance operation is to be consolidated to one primary server and one or more backup servers, plan for Ethernet cabling to interconnect equipment racks and ensure that all directors, switches, and server platforms have unique IP addresses.
- **Install equipment racks on a public LAN** — If a public LAN segment is to be used, determine from the customer's network administrator how to integrate the products and HAFM appliance. Ensure all access, security, and IP addressing issues are resolved.

Note: HP recommends that directors, switches, and the HAFM appliance be installed in a secure physical network domain to optimize security and avoid traffic problems.

- **Install remote user workstations** — Plan for access to the LAN segment containing the HAFM appliance if remote user workstations are required.

Task 6: Plan Network Addresses

Depending on the configuration of the LAN on which directors, switches, and the HAFM appliance are installed, plan network addressing as follows:

- If installing products and the HAFM appliance on a dedicated (private) LAN segment, there is no requirement to change any default network addresses. If multiple equipment racks are connected, ensure that all directors, switches, and servers have unique IP addresses. If new IP addresses are required, consult with the customer's network administrator.
- If installing products and the HAFM appliance on a public LAN containing other devices, default network addresses may require change to avoid address conflicts with existing devices.

  For the Director 2/64, Edge Switch 2/16, and Edge Switch 2/32, change the IP address, gateway address, and subnet mask through a remote terminal connected to the product's maintenance port.

  For the HAFM appliance, change the default network addresses through the **TCP/IP Properties** dialog box in Windows. In addition, assign and record a unique domain name service (DNS) name for the HAFM appliance and each director and switch.
- Gateway addresses may need to be configured for directors, switches, and the HAFM appliance if these devices connect to the LAN through a router or other gateway device.

The Ethernet connections for directors, switches, and the HAFM appliance have the following network addresses:

- Directors and switches:
  - Media access control (MAC) address is unique for each product. The MAC address is in ***xx.xx.xx.xx.xx.xx*** format, where each ***xx*** is a hexadecimal pair.
  - Factory preset and default IP address is **10.1.1.10**. If the Reset Configuration option is selected from the *Element Manager* application, the director or switch resets to this address.
  - Subnet mask is **255.0.0.0**.
  - Gateway address is **0.0.0.0**.
- HAFM appliance:
  - MAC address is unique.
  - IP address of the Ethernet adapter is **10.1.1.1**.

— Subnet mask is **255.0.0.0**.

— Gateway address is blank.

Task 7: Plan SNMP Support (Optional)

As an option, network administrators can use the *HAFM* application to configure an SNMP agent that runs on the HAFM appliance. This agent can be configured to send generic SNMP trap messages to up to 12 SNMP management workstations.

Administrators can also use the *Element Manager* application to configure an SNMP agent that runs on each director or switch. This agent can be configured to send generic SNMP trap messages to up to six SNMP management workstations.

Trap recipients can also access SNMP management information and may be granted permission to modify SNMP variables as follows:

- Assign and record product names, contact persons, descriptions, and locations to configure the products for SNMP management station access.
- Plan access to the director or switch LAN segment. This segment must connect to the LAN on which SNMP management workstations are installed.
- Obtain IP addresses and SNMP community names for management workstations that have access to products.
- Determine which (if any) management workstations can have write permission for SNMP variables.
- Obtain product-specific trap information from HP to load onto SNMP management workstations.

For additional information on SNMP, refer to the *HP StorageWorks SNMP Reference Guide for Edge Switches and Directors*.

Task 8: Plan E-Mail Notification (Optional)

As an option, network administrators can configure director and switch e-mail support. The following support considerations are required if the e-mail notification feature is used:

- Determine if e-mail notification is to be configured and used for significant system events.
- Determine which persons (up to five) require e-mail notification of significant director or switch events and record their e-mail addresses.
- Identify an attached e-mail server that supports the simple mail transfer protocol (SMTP) standard as defined in RFC 821.

Task 9: Establish Product and HAFM Appliance Security Measures

Effective network security measures are recommended for directors, switches, and the HAFM appliance. Physical access to the network should be limited and monitored, and password control should be strictly enforced.

When planning security measures, consider the following:

- Directors, switches, and the HAFM appliance are installed on a LAN segment and can be accessed by attached devices (including devices connected through a remote LAN). Access from remote devices is limited by installing the HAFM appliance and managed products in a secure physical network domain. HP recommends this approach.
- Access to products is possible through the maintenance port. This connection is for use by authorized service personnel only and should be carefully monitored.
- The number of remote workstations with access to the HAFM appliance and managed products can and should be restricted. Obtain IP addresses for workstations that should have exclusive access. Ensure that adequate security measures are established for the configured workstations.
- Carefully manage users (up to 16) who have access to the *HAFM* and *Element Manager* applications, and assign user names, passwords, and user rights.
- Ensure that adequate security controls are established for remote access software, including the Embedded Web Server.

Task 10: Plan Phone Connections

Plan for one or more telephone connections near the HAFM appliance for service personnel use. While performing a diagnostic or repair action, a service representative or network administrator at the HAFM appliance may require voice technical support through a telephone connection.

Task 11: Diagram the Planned Configuration

Determine peripheral devices that will connect to each director or switch, and determine whether and where connectivity should be limited (zoning). These devices may include servers, storage control devices, and other fabric elements in a multi-switch fabric.

Part of this task may have been performed when the configuration was determined. It might be helpful to draw the configuration diagram. Indicate distances in the diagram if necessary. Transfer information from the configuration diagram to the product planning worksheet provided as part of "Task 13: Complete the Planning Worksheet" on page 191.

Task 12: Assign Port Names and Nicknames

During the planning process, consider assigning names to director and switch ports based upon devices connected to the ports. Though not required, port naming provides convenience and ease of use. Port naming also documents devices that connect through individual ports and identifies what is attached to each port. When it is necessary to change port connectivity, port names make it easier to identify the ports and attached end devices.

Also consider assigning nicknames to device and fabric World Wide Names (WWNs). Though not required, nicknaming provides a useful substitute for the cryptic eight-byte WWN. Once a nickname is assigned, it is referenced throughout the *HAFM* application.

Transfer port names and nicknames to the product planning worksheet provided as part of "Task 13: Complete the Planning Worksheet" on page 191.

Rules for Port Names

Port names can be up to 24 alphanumeric characters long. Spaces, hyphens (**-**), and underscores (_) are allowed within the name. Each port name must be unique for a director; however, the same port name can be used on separate directors and switches. HP recommends that unique port names be used, particularly within a complex multi-switch fabric. Example port names include:

```
Lab server
Test system-2
Printer_001
```

Rules for Nicknames

Nicknames can be up to 32 alphanumeric characters long. Spaces, hyphens (**-**), and underscores (_) are allowed within the name. Each nickname must be unique (corresponding to a unique WWN). Example nicknames include:

```
Fabric-1
Host system
DASD_001
```

Task 13: Complete the Planning Worksheet

The planning worksheet included in this task is a four-page form that depicts port assignments for a director or switch. The worksheet lists 64 ports and provides fields to identify devices that connect to the ports.

Transfer information from the configuration diagram (completed while performing "Task 11: Diagram the Planned Configuration" on page 189) to the worksheet, and transfer port names and nicknames (assigned while performing "Task 12: Assign Port Names and Nicknames" on page 190). In addition, indicate all unused ports. Retain the planning worksheet as part of a permanent record.

Table 6: Product Planning Worksheet (Page 1 of 4)

Director or Switch Name: ______ IP Address: ______ Unit Name: ______			Attached Devices			
Port	Port Name	Location	Type	Model	IP Address	Zone
00						
01						
02						
03						
04						
05						
06						
07						
08						
09						
10						
11						
12						
13						
14						
15						

Product Planning Worksheet (Page 2 of 4)

Director or Switch Name: ______ IP Address: ______ Unit Name: ______			Attached Devices			
Port	Port Name	Location	Type	Model	IP Address	Zone
16						
17						
18						
19						
20						
21						
22						
23						
24						
25						
26						
27						
28						
29						
30						
31						

Product Planning Worksheet (Page 3 of 4)

Director or Switch Name: ______ IP Address: ______ Unit Name: ______			Attached Devices			
Port	Port Name	Location	Type	Model	IP Address	Zone
32						

Product Planning Worksheet (Page 3 of 4)

33						
34						
35						
36						
37						
38						
39						
40						
41						
42						
43						
44						
45						
46						
47						

Product Planning Worksheet (Page 4 of 4)

Director or Switch Name: ____ IP Address: ____ Unit Name: ____			Attached Devices			
Port	Port Name	Location	Type	Model	IP Address	Zone
48						
49						
50						
51						
52						
53						
54						

Product Planning Worksheet (Page 4 of 4)

55						
56						
57						
58						
59						
60						
61						
62						
63						

Task 14: Plan AC Power

Plan for facility power sources for each equipment rack. Directors and switches in the rack operate at 50 to 60 Hertz (Hz) and 100 to 240 volts alternating current (VAC) and require a minimum dedicated 5-ampere service. If two power sources are supplied (optional but recommended for high availability), the equipment rack contains two customer-specified external power cords. Each cord should be connected to a separate power circuit or both should be connected to an uninterruptable power supply (UPS). Several types of power cables and plugs are available to meet local electrical requirements.

WARNING: An insulated grounding conductor identical in size, insulating material, and thickness to the grounded and ungrounded branch-circuit supply conductors (except it is green, with or without one or more yellow stripes) shall be installed as part of the branch circuit supplying the product. The grounding conductor described shall be connected to ground at the product, or if supplied by a separately derived system at the supply transformer or motor generator. The plug receptacles near the product shall all be a grounding type, and grounding conductors serving these receptacles shall be connected to ground at the product.

Keep all power cables out of high-traffic areas for safety and to avoid power interruption caused by accidentally unplugging the product or equipment rack.

Task 15: Plan a Multi-Switch Fabric (Optional)

If a multi-switch fabric topology is to be implemented, carefully plan the physical characteristics and performance objectives of the topology. Include the proposed number of fabric elements, characteristics of attached devices, cost, nondisruptive growth requirements, and service requirements.

When two or more fabric elements are connected through ISLs to form a fabric, the elements must have compatible operating parameters, compatible name server zoning configurations, and unique domain identifications (IDs). Planning for a fabric must be carefully coordinated with planning for zoned configurations.

Consider the following factors when planning for a multi-switch fabric:

- **Fabric topology limits** — Consider the practical number of fabric elements (theoretical maximum of 31, practical limit of 24), number of ISLs per element, hop count (maximum of 3), and distance limitations (limited by port type and cable availability).
- **Bandwidth** — Consider using multiple ISLs to increase the total bandwidth available between two fabric elements.
- **Load balancing** — If heavy traffic between devices is expected, consider installing multiple ISLs to create multiple minimum-hop paths for load balancing.
- **Principal switch selection** — If required, plan which fabric element is to be assigned principal switch duties for the fabric.
- **Critical operations** — Consider routing paths that transfer data for critical operations directly through one director or switch and not through the fabric.

Planning and implementing a multi-switch fabric is a complex and difficult task. HP recommends that you obtain planning assistance from our professional services organization before implementing a fabric topology.

Task 16: Plan Zone Sets for Multiple Products (Optional)

If name server zoning is to be implemented, carefully plan the characteristics and security objectives (differentiation of operating systems, data sets, user groups, devices, or processes) of zone members, zones, and zone sets.

If a fabric topology is implemented, zoning is configured on a fabric-wide basis. Planning for zoned configurations must be carefully coordinated with planning a fabric topology.

Consider the following factors when planning to implement name server zoning:

- **Zone and zone set naming conventions** — Directors and switches conform to the open fabric naming convention by using the following zone and zone set naming rules:
  - Zone and zone set names can be up to 64 characters long.
  - The first character of the name must be an upper-case alpha character (**A** through **Z**) or lower-case alpha character (**a** through **z**).
  - Characters other than the first character can be upper-case or lower-case alphanumeric characters (**A** through **Z**, **a** through **z**, or **0** through **9**), a dollar sign (**$**), hyphen (**-**), caret (**^**), or underscore (_).
- **Zone members specified by port number or WWN** — Consider if zoning is to be implemented by port number or WWN. Because changes to port connections or fiber-optic cable configurations may disrupt zone operation, zoning by WWN is recommended.
- **Zoning implications for a multi-switch fabric** — To ensure zoning is consistent across a multi-switch fabric, directors and switches must have compatible operating parameters and unique domain IDs, the active zone set name must be consistent, and zones with the same name must have identical elements.
- **Server and storage device access control** — In addition to zoning, consider implementing server-level access control (persistent binding) and storage-level access control.

Planning and implementing zones and zone sets is a complex and difficult task, especially for multi-switch fabrics. HP recommends that you obtain planning assistance from our professional services organization before implementing a director or switch zoning feature.

Index

H

I

J

L

M